



AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in the Pipe You Own

Your category is attacked from above and from below at the same time. Your twenty years of proprietary project files are training data for the moat you have not yet shipped. Your scorecard measures platform uptime and feature velocity, and neither is the 2026 question. A briefing for CTOs, CPOs, and CPTOs at mid-market European specialist research and intelligence firms.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market specialist
research and intelligence firms

DATE

April 2026

Tuesday 10:45, client QBR

Every CTO at a mid-sized European research firm has had the same Tuesday. Your research director runs the quarterly pipeline with your biggest FS subscriber. Their research director mentions, casually, that they have been running your reports through Claude on their own data and can reproduce about seventy percent of the analytical depth in-house at one-tenth the subscription cost. He is not angry. He is asking what you deliver in the other thirty percent. Your subscription contract is up for renewal in November. Your CFO has a PE sponsor on a call Friday. Your VP Engineering's last senior data scientist left for a Perplexity competitor three weeks ago.

You have run this platform for ten years. The pattern used to be simple. Build the research pipeline. Hold the editorial bar. Ship the subscription portal. Meet the board on platform uptime and feature shipping. The business was understandable.

Something changed. Your category is attacked from above and below at the same time. Above, your enterprise clients with Claude Enterprise and ChatGPT Enterprise access are producing the analysis they used to buy. Below, Perplexity Enterprise, Elicit, Glass AI, and Consensus are shipping synthesis at a fraction of your unit economics on public sources. What remains between the two attacks is the thirty percent your senior analysts produce on top of twenty years of proprietary project files and a forecast-accuracy track record nobody else has.

You are not running one R&D function. You are running two, and only one of them is on your scorecard. One funds the subscription platform and the client portal. The other funds what has to exist by 2028: the indexed corpus of your project files, the data rights you have not renegotiated, the forecast-accuracy telemetry you have not wired, the retrieval architecture that makes your twenty-year pipe queryable by your analysts and your clients. 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 R&D function you run, and the three questions your team should answer before the next board cycle.

Your category is attacked from above and below at once

The attack from above is the one every research CEO now reads in the renewal meeting. Your enterprise clients have access to Claude Enterprise, ChatGPT Enterprise, and a data team that did not exist three years ago. They are producing seventy percent of the analysis they used to buy. The quality is imperfect but it is good enough to reset their expectations permanently. The pricing conversation is not "your subscription is too expensive." It is "we can do most of this in-house now, what are you delivering that we cannot?" That conversation ends with a renewal at reduced scope, or a churn, or a migration to a per-query model the subscription P&L cannot absorb.

The attack from below is the one the AI-natives are running with different economics. Perplexity Enterprise, Elicit, Glass AI, and Consensus each own a slice of what used to be a research-firm category. They ship synthesis at one-tenth of your unit economics on public sources and they are improving weekly. Your client compares your quote side by side with theirs. They are not asking which research is better. They are asking which one their CFO can defend against the board. The middle that remains, between the client doing it themselves and the AI-native shipping at cost, is the thirty percent your senior analysts deliver on top of your proprietary pipe and your accumulated forecast track record. That middle is defensible. It is also the only thing that is defensible.

Your subscription product commoditises every eighteen months as the next model rolls. Your indexed corpus of twenty years of project files is the asset the AI-native cannot replicate without your client base.

The category's speed mismatch is the third force. AI-natives build from scratch with no legacy, no fifteen-year-old research platform, no client portal that has not been touched since 2022. Your firm carries all three. You cannot pivot at their speed on their terms. What you can do is identify what is genuinely defensible in your product, stop defending the pieces that are not, and ship the corpus-and-retrieval architecture they cannot build. The firms that are already moving are quiet about it. The ones shipping "AI-augmented research" as a feature in the subscription portal are the ones the board will replace in 2028.

Your R&D function is feeling it every day

Your build engine. You have forty engineers, a fifteen-year-old research platform, a client portal nobody has touched since 2022, and an analyst workbench built as a Sharepoint replacement. The AI-native shipping into your category has twelve engineers and a six-month-old stack. They ship weekly. You ship quarterly. The backlog is dominated by integration debt, CMS replatforming, and data-pipeline remediation. Copilot adoption helps at the margin and does not move the number on the work that matters.

Your platform and data. Your twenty years of historical project files sit in a CMS that was designed for distribution, not retrieval. Your research corpus is not indexed. Your embeddings are last generation or do not exist. Your data rights across the top twenty client contracts were written before anyone thought a model would train on them. The moat you can ship is an indexed, retrieval-ready, model-agnostic corpus over that pipe, with data rights renegotiated across the next twenty-four months of renewals. That stack compounds. A two-person AI-native cannot replicate it because they do not have your client base or your back catalogue.

Your product thesis. Your board asked for the firm's three-sentence product thesis last quarter. Your answer was the one from three years ago. Your rate card prices access. Your clients are paying for access to reports they can increasingly generate themselves. The firms that survive price insight outcomes, not access. That is a capital-allocation question about where the moat lives, not a pricing one. Your CFO and your Head of Research need a joint document with you before the next renewal cycle, not after.

Your R&D bench. You stopped hiring junior analysts in 2024 because the routine research tasks were moving to agents. The junior pipeline that used to train juniors into mid-tier analysts has collapsed. Two of your three most tenured analysts had a recruiter call from a Perplexity competitor this month. Your five-year senior bench is empty. OpenAI and Anthropic resumed junior hiring in late 2025 to pair juniors with seniors on judgment-intensive work. The firm that cuts for this year's number has a cheaper function in Q4 and an empty bench in 2031.

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 shortlist for a Chief Data Officer. Every month without a written boundary, the AI thesis drifts further from the seat that can actually defend it. The Q1 boundary agreement lands before the shortlist becomes a hire, not after.

The middle is the moat

Your sector is being repriced from above because enterprise clients now produce what they used to buy. Your sector is also being attacked from below because AI-natives ship synthesis at a fraction of your unit economics on public sources. Same force, two altitudes. But there is a third thing happening that most mid-market research CTOs have not named yet.

The moat that matters in 2028 is neither the subscription product nor the AI feature bolted on top of it. It is the pipe underneath both. Your twenty years of project files, your methodology records, your forecast-accuracy telemetry, and the integration depth with your top twenty clients. The product commoditises. The pipe compounds. An AI-native with a better retrieval architecture cannot ship that pipe because they do not have your client base, your back catalogue, or your analyst editorial history. You do. The window closes in roughly eighteen months, which is the time between now and when the AI-native has enough of their own client footprint to start building a comparable corpus from scratch.

In 1985, a news agency faced the same two-front attack. Wire-service rivals were cutting production cost with new telecoms. Banks with internal research teams were reproducing the wire's market commentary in-house. The agency's CTO did not compete on faster filing or smarter reports. He built the terminal. The terminal was not a news product. It was an indexed, queryable corpus of market data, news, editorial metadata, and historical pricing, priced on a subscription the bank analyst paid for because no competitor could replicate the back catalogue. By 1995, the news-wire business had commoditised. The terminal business had not. Same shape, different year.

You can tell which camp a firm is in from a single question. Is the R&D budget you defended at the last board meeting funding the subscription product, or funding the indexed pipe underneath it? Most mid-market research firms cannot answer. The ones that can have already made the shift.

You have twelve to eighteen months. The CTOs who claim authorship of the pipe in Q1 author the decade. The ones who wait until a Chief Data Officer arrives inherit a seat that has narrowed to platform uptime and portal UX, and a scorecard that stopped mattering the quarter the biggest subscriber brought the analysis in-house.

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

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

THE DEFAULT PATH

Ship AI features into the subscription portal

Adds a chat interface on top of the existing reports. Features commoditise. The moat does not move.

Run one R&D budget, blended hurdle rate

Portal bug fixes and corpus indexing compete on the same quarter. Portal wins every time. The moat never ships.

Price on access to reports

Clients renew at reduced scope or churn. AI-natives win net-new on unit economics you cannot match.

Cut the junior analyst layer for the Q1 number

Cheaper function this quarter. Empty senior bench by 2031. External hires at premiums the firm cannot absorb.

THE WINNERS' PATH

Index the pipe underneath the product

Twenty years of project files become a retrieval-ready corpus. Model-agnostic. The AI-native cannot build it.

Split the budget in two, two hurdle rates

Plateau compresses production cost. Compounding funds the pipe, the data rights, the forecast telemetry.

Price on outcome, on protected P&L

Insight advisory priced on forecast accuracy. Subscription holds installed base. Outcome line wins net-new.

Pair juniors with seniors and agents on insight work

Juniors start on audit and paired synthesis, not routine production. Eighteen-month apprenticeship on judgment.

Same engineering team. Same twelve months. Opposite outcomes. The difference is whether the CTO authors the pipe or signs the requisition for a subscription-portal refresh.

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

Three questions every research 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 your engineering team shipping the platform or fighting integration debt?

Forty engineers, a fifteen-year-old platform, a client portal untouched since 2022, an analyst workbench built on Sharepoint. The AI-native has twelve engineers and a six-month-old stack. Copilot adoption moves cycle time at the margin. The real bottleneck is the backlog dominated by CMS debt and pipeline remediation. Buy what the tool layer carries. Retire the debt on a fixed schedule. Route the freed hours into the pipe.

02 · PRODUCT DEFENSIBILITY

What does your product do that a client with Claude and an AI-native at ten percent cost cannot copy?

Above, your clients do seventy percent internally. Below, AI-natives ship synthesis at one-tenth your unit economics on public sources. What remains is the thirty percent your senior analysts deliver on top of twenty years of proprietary project files and forecast-accuracy telemetry. That middle is defensible only if it is encoded. Today it sits in PDFs and three analysts' heads, two of whom had recruiter calls last month.

03 · R&D CAPITAL ALLOCATION

Is your R&D budget one instrument or two?

Plateau capital runs the subscription platform at lower cost. Compounding capital builds the indexed corpus, the retrieval stack, the data-rights renegotiation, and the forecast telemetry. On one hurdle rate the first wins every quarter. On one scorecard the second does not exist. The CTO who walks in with two budgets, each defended separately, is the one the board trusts with the capital for the next three years.

What the firms winning this play are actually doing

Four concrete moves any research CTO can start this quarter. Each addresses one sub-function. Pulled together, they become one architecture the firm builds on.

1 **Build engine: retire integration debt on a fixed schedule, free the team for the pipe**

Instrument cycle time and review depth, but fight the real bottleneck first. Audit the top ten integration debts blocking platform velocity. Schedule them in a rolling ninety-day cadence. Route agent-assisted engineering into the remediation work, not into feature shipping. Free the best four engineers for the pipe programme. The dashboard is no longer feature velocity. It is engineer-hours redirected to the corpus, the retrieval stack, and the telemetry.

2 **Platform and data: ship the indexed pipe underneath the subscription product**

Three programmes that stack over eighteen months. A retrieval-ready index across twenty years of historical project files, with embedding strategy portable across GPT-5, Claude 5, and Gemini 3. A data-rights audit across the top twenty client contracts, with renegotiation plans synchronised to each renewal window. A forecast-accuracy telemetry pipeline tying every published prediction to its realised outcome, becoming both the eval harness for your agents and the public differentiator against AI-native synthesis.

3 **Product thesis: stand one insight-advisory line on protected P&L**

Do not migrate the whole book. Keep subscription for the installed base that is not ready to reprice. Ship insight advisory on a separate P&L, priced on outcome rather than access. Outcome shapes: forecast-accuracy guarantees, decision-impact clauses, or named-deliverable pricing keyed to the client's own success metric. Write the joint Pricing Architecture Document with your CFO and Head of Research in Q1, before the first outcome deal closes. The document makes the transition auditable inside your firm and explicit for the buyer.

4 **R&D bench: rebuild the junior analyst pathway around senior and agent pairing**

Your 2024 decision to stop hiring junior analysts saves Q1. It empties the 2031 senior bench. Reverse it. Juniors in 2026 do not start on routine research production. The agent layer absorbed that work. They start on audit, exception handling, and paired synthesis with a named senior for their first eighteen months. The career curve compresses from five years to months because the senior's judgment becomes legible in the agent-and-junior review loop. By 2029 you have a mid-tier bench the competition cannot hire away because it did not exist anywhere else to hire from.

The only firm that joins R&D authorship to working infrastructure

The research CTO has four places to go. None can do what we do.

The big tech-transformation practices ship a standardised AI-governance framework scaled across a sector book. Their cost structure requires reusability. Your seat is idiosyncratic. A reusable framework on an idiosyncratic seat produces governance theatre you cannot defend to the board.

The hyperscaler professional services arms price the engagement as enablement for cloud consumption. Their business model routes every build-or-buy decision through their pricing for three years. The proprietary corpus you send through their API is training data for a general model.

The AI-platform vendors have a quote on your desk. Perplexity Enterprise, Anthropic enterprise, OpenAI enterprise, research-synthesis platforms. Some are selling you tooling. Some are your existential competitive threat framed as a partner. Telling you to encode your corpus so their synthesis cannot reach it is structurally impossible for them to say.

An in-house or fractional Chief Data Officer splits the seat. Strategy-side CDO is revenue-measured; platform-side CDO is uptime-measured. One person cannot hold both. Two-thirds of newly-appointed CAIOs and CDOs spend year one compiling inventories instead of setting direction.

The four options you have already looked at. None can run the pipe conversation the seat requires. GRAIL is the fifth.

WHAT YOU GET

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

Working infrastructure your team uses the same week

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

The start of a data-and-retrieval moat your platform 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 R&D leadership. Working infrastructure from week one.

We start with your R&D 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 R&D function. Build the first piece of infrastructure in the room: a capital-allocation agent that takes your real R&D budget, splits the plateau work from the compounding work, 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 R&D infrastructure

A half-day build. Your VP Engineering, Head of Data, Head of Research, and Head of Product each leave with five to ten pieces of working infrastructure tuned to their function. Corpus-indexing prototype. Retrieval-quality evaluator. Data-rights renegotiation playbook. Forecast-accuracy telemetry scaffold. 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 budget 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 Head of Research. The eighteen-month pipe build plan.

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

For your next R&D review

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

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

Name the AI-native in your category and the percent of your unit economics they quote at. If you cannot name the top three, that is the first thing the briefing changes.

If your two most tenured analysts left for a Perplexity competitor tomorrow, how many months to reconstruct their pattern-recognition from files the firm already has?

Your biggest subscriber runs your reports through Claude on their own data. Where does your product deliver value they cannot self-serve, and is that value priced on outcome or on access?

Heidrick's February 2026 data: thirty percent CIO, twenty-three percent CDAO, twenty-one percent CAIO on the AI-strategy-ownership list. Zero percent CTO-CPO. Is your Q1 boundary agreement with the CEO written, or is it waiting until after the CDO shortlist lands?

If any of these produces a disagreement in your R&D 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