



AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in the Moat You Build

Your R&D budget is two instruments pretending to be one. Your product sits above a model layer that commoditises every eighteen months. Your scorecard measures feature velocity and engineering headcount, and neither is the 2026 question. A briefing for CTOs, CPOs, and CPTOs at mid-market European B2B SaaS and tech firms.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market B2B SaaS
technology and product leaders

DATE

April 2026

Monday 9:15, R&D review

Every CTO at a mid-sized European B2B SaaS firm has had the same Monday. Sprint velocity up twenty-two percent quarter on quarter. Copilot adoption at seventy-one percent. Cycle time from merge to production flat against last quarter. Your VP Engineering opens the review with the dashboard. He is mid-sentence when your phone buzzes. It is the CTO at your biggest customer. One sentence: "We've started a Sierra pilot on the support-resolution module, happy to talk." Your VP Engineering mentioned a recruiter call from Cursor over coffee Thursday. Your CEO has a CAIO shortlist from the board on his desk. Your CFO forwarded the Shopify R&D line this morning: "Worth discussing before Q2 planning."

You have run this function for ten years. The pattern used to be simple. Hire strong engineers. Hold the review bar. Ship the roadmap. Meet the board on feature velocity and engineering headcount. The business was understandable.

Something changed. What used to take fifty engineers now runs on five. That compression arrived faster than the discipline to absorb it. Ninety-five percent of corporate AI deployments produce no measurable P&L 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 AI-native at ten percent of your unit economics just landed its first named reference in your category. Your VP Engineering carries architectural reasoning that has never been written down, and two Sierra 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 ship. The other funds what has to exist by 2028: the data rights you have not renegotiated, the agent-memory store you have not built, the integration depth that outlasts three model swaps, the architectural reasoning sitting in three engineers' heads. 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.

B2B SaaS just lost the equation that priced the category

The AI-natives in your category are not experimenting. They are building. Sierra shipped to a hundred million in ARR in twenty-one months on outcomes priced at roughly ten percent of your per-seat unit economics. Harvey hit forty-two percent of the AmLaw 100 in one year. Decagon, Legora, Peec AI, Cursor each own a slice of what you thought was a feature moat. Your customers' procurement teams run comparison quotes where your per-seat number sits next to the AI-native's per-outcome number. They are not asking which is the better product. They are asking which pricing model their CFO can defend.

The moat your product ran on for a decade has moved. The feature layer commoditises fast now. A capability that used to take a 150-person team three years to ship can be rebuilt by a 20-person AI-native in twelve to eighteen months, on a modern architecture, priced on an outcome your customer can measure. Your feature edge was never going to hold. What can hold is the layer underneath the model: the accumulated understanding of your customers, the integration depth your product has wired into their systems, the data rights that let your agents be grounded in their reality. Those compound. They take years to build and they cannot be bought. An AI-native cannot ship them because they do not have your customer base.

Your feature shipped on GPT-5 is on an eighteen-month obsolescence curve. Your integration depth with your top ten customers is the asset the AI-native cannot replicate without your customer base.

The industry's speed mismatch is the third force. AI-natives build from scratch with no legacy, no eighteen-month roadmap cycles, no complex procurement to answer to inside their own house. Your firm has 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 what is not, and ship the moat they cannot build. The firms that are already moving are quiet about it. The ones that are shipping Copilot and calling it an AI strategy are the ones the board will replace in 2028.

Your R&D function 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 people are pushing ten-thousand-line PRs barely reviewed. The foundational layer rots. Twelve months in, your team cannot execute a category rewrite when one arrives. The review discipline that worked in 2022 is training the wrong practice on agent-generated volume now.

Your platform and data. Spotify's best developers have not written code since December because Spotify spent four years on the platform foundation. 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: data rights renegotiated across your next twenty-four months of contract renewals, a model-agnostic agent-memory store portable across model swaps, an integration depth graph published for your top ten customer systems. The model layer commoditises. The data layer compounds.

Your product thesis. Your board asked for the firm's three-sentence thesis last quarter. Your answer was the one from three years ago. Your rate card assumes AI features will price like legacy features. They do not. Sierra runs three pricing mechanics. You run one. That is a capital-allocation question about where the moat gets built, not a pricing one. Your CFO and commercial leader need a joint document with you before the AI-native forces the conversation on their terms.

Your R&D bench. You stopped hiring juniors in 2024. The routine work that trained juniors into seniors is now running on agents. No routine work, no training ground, no senior bench in 2030, no internal CPTO 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 for this year's number has a cheaper function in Q4 and an empty bench in five years.

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 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 above the model layer

Your sector is repricing from below because AI-natives are shipping outcomes-priced products at ten percent of your unit economics. Your function is compressing because five engineers 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. There is a third thing happening that most mid-market B2B SaaS CTOs have not named yet.

The moat that matters in 2028 is not above the model layer. It is underneath it. The product feature you ship this quarter on GPT-5 is on an eighteen-month obsolescence curve. The next foundation model rolls in 2027 and your feature commoditises. What does not commoditise: the data you have rights to, the agent-memory corpus you own, and the integration depth you have published for your top ten customer systems. A two-person AI-native cannot build that layer because they do not have the customer base that generates the signal. 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 customer footprint to start building the same layer.

In 2006, a handset manufacturer shipped the best-engineered product in the world. The queue for their next flagship was eighteen months out. Their CTO's scorecard was battery life, keyboard quality, enterprise email reliability, carrier relationships. Every metric was green. Apple shipped a platform. By 2011 that handset manufacturer's market share in its own category had fallen from forty-three percent to ten. The 2006 scorecard never turned red. It just stopped being the one the market was scoring. The job had not changed. The unit of judgment had.

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 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 claim the authorship of this decision in Q1 author the decade. The ones who wait until the CAIO hire arrives inherit a seat that has narrowed to platform uptime and feature velocity, and a scorecard that stopped mattering the quarter their biggest account opened a Sierra pilot.

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.
Ten-thousand-line PRs barely reviewed.
Foundation rots under a green dashboard.

Run AI budget through IT, sign the requisition

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

Ship AI features free into the existing rate card

Customer value rises thirty to sixty percent. ARPU flat. Sierra lands your first reference loss at ten percent of your unit economics.

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 firm cannot absorb.

THE WINNERS' PATH

Buy the tool layer, redesign the review bar

Review-time floor per line. ADRs mandatory on agent-generated changes. Eval harness as first-class infrastructure.

Route R&D capital through the M&A hurdle

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

Stand one outcome-priced line on protected P&L

Parallel pricing architecture with CFO and commercial. Seat-based holds installed base. Outcome-based wins net-new. No cannibalisation.

Pair juniors with seniors on compounding work

OpenAI and Anthropic resumed junior hiring in late 2025. Eighteen-month apprenticeship on judgment, not five years on routine.

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

THREE QUESTIONS

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

Three questions every B2B SaaS CTO is tracking. The third is the crux. The first two are how you earn the right to answer it.

01 · BUILD VELOCITY

Is our engineering speed shipping production-grade output, or demos that fall over?

Copilot adoption at seventy percent. Cycle time flat. The anonymous engineering manager is in your firm too: ten-thousand-line PRs barely reviewed, foundation rotting. Spotify spent four years building 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 product do that a Sierra at ten percent of our unit economics cannot copy?

Your feature edge commoditises every eighteen months when the next foundation model rolls. Your integration depth and data rights compound. Your architectural reasoning sits in three engineers' heads, and two of them had recruiter calls from an AI-native last month. The moat sits underneath the model layer, not above it. The window to build it is eighteen months.

03 · R&D CAPITAL ALLOCATION

Is our R&D budget one instrument or two?

Plateau capital runs the existing product at lower unit cost. Compounding capital builds the 2028 moat. 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 same 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 B2B SaaS 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: golden set, rubric, automated grading, no harness, no ship. Rebalance the senior-to-junior ratio around the composite rather than preserving the 2022 team shape.

2**Platform and data: build the moat underneath the model layer**

Three concrete programmes that stack over eighteen months. Data rights audit across your top customer contracts, with a renegotiation plan tied to your next twenty-four months of renewals. A model-agnostic agent-memory store portable across GPT-5, Claude 5, and Gemini 3, separated from any specific model provider's context window. An integration depth graph published for your top ten customer systems, maintained as structured metadata your sales team can walk into a procurement conversation with.

3**Product thesis: stand one outcome-priced line on protected P&L**

Do not migrate the whole book. Keep seat-based for the installed base that does not want to switch. Ship outcome-priced on a separate P&L for net-new and for customers explicitly asking. Write the cannibalisation trigger and kill criteria in Q1: if outcome revenue on an account exceeds seat revenue, trigger the migration conversation. The product architecture supports both mechanics natively. The joint Pricing Architecture Document with your CFO and commercial leader is signed before the first outcome-priced deal closes.

4**R&D 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 audit, exception-handling, and paired composite work with a named senior for their first eighteen months. The career curve compresses from five years to 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 competition cannot hire away.

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

The B2B SaaS 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 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, agent-workflow platforms. Their product is a build-velocity intervention measured on seats. Telling you to spend freed capacity on the data moat instead of expanded seats is structurally impossible for them to say.

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 of newly-appointed CAIOs spend year one compiling inventories instead of setting direction.

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

WHAT YOU GET

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

Working infrastructure your team uses the same week

Custom agents built on your code, your customers, your stack

The start of a data-and-integration 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 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 R&D infrastructure

A half-day build. Your VP Engineering, Head of Product, Head of Platform, and Head of Data each leave with five to ten pieces of working infrastructure tuned to their function. Review-bar audit agent. Architectural-decision scaffolding. Data-rights renegotiation playbook. Integration depth graph prototype. 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 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 commercial leader. The eighteen-month moat 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 splits plateau from compounding, and what is the kill criterion on each?

Name the AI-native in your category. What is their pricing mechanic, and at what percent of your unit economics are they quoting? If you cannot name the top three, that is the first thing the briefing changes.

If your VP Engineering or two founding engineers left for Sierra or Cursor tomorrow, how many months would it take to reconstruct the architectural reasoning behind your core product decisions?

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 foundation layer trained on?

Heidrick's February 2026 data: thirty percent CIO, twenty-three percent CDAO, twenty-one percent CAIO, one percent COO, zero percent CTO-CPO on the 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 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