

AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in the Telemetry You Own

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 platform uptime and feature velocity, and neither is the 2026 question. A briefing for CTOs, CPOs, and Heads of Platform at mid-market European energy and environment firms.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market energy and environment technology leaders

DATE

April 2026

Monday 7:40, forecast review

Every CTO at a mid-sized European energy and environment firm has had the same Monday. Your forecast-error dashboard shows balancing-interval error at one point eight percent, point-four wider than last week. Your duty controls engineer has left a note. Two wind assets in the north cluster had a ramp-rate event Sunday the model did not predict. Your phone buzzes. It is your biggest TSO customer. One sentence: "Our board is looking at an AI-native pilot for forecast-error reduction, happy to talk before we shortlist." Your senior data scientist mentioned a recruiter call from a developer over coffee Thursday. Your CEO has a CAIO shortlist from the board on his desk. Your CFO forwarded the Shopify R&D trajectory this morning: "Worth discussing before Q2 planning."

You have run this function for ten years. The pattern used to be simple. Hire strong engineers. Ship forecast, asset-intelligence, or diligence code at the release cadence. Hold the integration bar against the customer's SCADA environment. Meet the board on platform uptime and feature velocity. The business was understandable.

Something changed. What used to take fifty engineers on a forecasting or asset-intelligence product now runs on five. 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 is beginning to pitch into asset-ops optimisation, forecast-error reduction, and diligence automation. Your senior controls engineer carries integration reasoning that has never been written down, and a developer client is offering forty percent more to hire him.

You are not running one R&D function. You are running two, and only one is on your scorecard. One funds what you already ship. The other funds what has to exist by 2028: the asset-telemetry rights you have not renegotiated, the regulatory-precedent corpus you have not encoded, the integration depth into the top ten customer SCADA environments that outlasts three model swaps. 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.

The knowledge you generate this year is worth more next year than today

Energy and environment is the one sector in Europe where the knowledge problem runs in reverse. Every other industry is trying to capture expertise before it retires. You are trying to encode expertise faster than you create it. Every offshore-wind assessment, every hydrogen pathway, every BESS permit, every grid-connection response, every forecast-model calibration against a Nordic winter operating profile generates asset and regulatory precedent worth more next year than today. The firm where that knowledge compounds institutionally has a structural advantage competitors cannot hire away. The firm where it lives in three senior heads loses the advantage the month two of them take a developer offer at forty percent premium.

The AI-natives in your category are starting to arrive. They are pitching into asset-ops optimisation, forecast-error reduction, trading-signal generation, diligence automation, and regulatory-submission drafting. Sierra, Harvey, Decagon, and the sector-specific entrants shipped to first reference customers in the past twelve months at roughly ten percent of your per-project or per-seat unit economics. Your TSO, DSO, and IPP customers' procurement teams run comparison quotes where your price 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.

Your feature shipped on the current foundation model is on an eighteen-month obsolescence curve. Your integration depth into the top ten customer SCADA environments is the asset the AI-native cannot replicate without your customer base.

Between 2015 and 2020 the sector watched a major industrial-software platform launch with enormous capital and a promise to unify asset intelligence across every vertical. It retreated. It tried to build a horizontal platform without owning vertical asset-data rights and without winning the integration fight into specific customer operating environments. The narrower competitors with tighter integration took the market. That memory is fresh. The CTOs who remember it and act on it are already shipping the data-rights audit and the integration-depth graph this quarter. The ones running Copilot across the engineering team 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 across your engineering team. Cycle time has not moved. Velocity is up on the dashboard. Your controls engineers are pushing ten-thousand-line PRs barely reviewed, and the integration code that touches your customer SCADA environments is where the foundation rots fastest. Twelve months in, your team cannot execute a category rewrite when one arrives. The review discipline that worked on a four-person controls team 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 platform foundation. You cannot replicate four years in eighteen months while also defending against an AI-native at ten percent of your unit economics. You can buy build-side discipline from the tool layer. Redirect compounding capital into the layer the tool vendors cannot sell: asset-telemetry rights renegotiated across your next twenty-four months of customer contract renewals, a model-agnostic agent-memory store for asset behaviour and regulatory precedent, an integration-depth graph published for your top ten customer SCADA and DCS environments. Model layer commoditises. This 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-augmented features will price like legacy features. They do not. The AI-natives in asset-ops optimisation, forecasting, and diligence are quoting per-outcome. You quote per-project or per-seat. 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. Your senior controls engineers, senior data scientists, and senior regulatory specialists carry the reasoning behind every integration, every forecast-model calibration, every permit strategy your product runs on. Two of them have recruiter calls from developers at forty percent salary premium this quarter. You stopped hiring juniors in 2024. No routine work, no training ground, no senior bench in 2030. 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 the integration depth you wrote off as debt

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. But there is a third thing happening that most mid-market energy and environment CTOs have not named yet.

The integration depth you have been writing off as technical debt is the moat you have already built. Twenty years of field-conditioned integration into customer SCADA, DCS, and CIM environments. Every permit response you filed. Every ramp-rate event your model learned from. Every grid-code nuance your regulatory specialist walked the TSO through. A two-person AI-native cannot replicate 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 first AI-native entrants have 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 TSO customer opened a pilot with someone else.

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 touching SCADA integration code, barely reviewed.

Route AI budget through IT, sign the requisition

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

Quote AI features into the existing project fee

Customer value rises thirty percent. Fee flat. First AI-native reference loss at a tenth of your unit economics.

Cut the junior layer for the Q1 headcount win

Cheaper function this quarter. Empty succession bench by 2031. External senior hire at a premium the firm cannot absorb.

THE WINNERS' PATH

Buy the tool layer, redesign the review bar

Review-time floor per line on production integrations. ADRs mandatory on agent 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

Joint Pricing Architecture with CFO and commercial. Project-fee holds installed base. Outcome-priced wins net-new.

Pair juniors with seniors on compounding work

Eighteen-month apprenticeship on field judgment, permit strategy, and integration depth. 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.

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

Three questions every energy and environment 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 production-grade integrations, or demos that break at the first grid-stress event?

Copilot adoption at seventy percent. Cycle time flat. Ten-thousand-line PRs touching SCADA integration code, barely reviewed. 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 into a live customer environment.

02 · PRODUCT DEFENSIBILITY

What does our product do that an AI-native at ten percent of our unit economics cannot copy?

Your feature edge commoditises every eighteen months when the next foundation model rolls. Your asset-telemetry rights, your integration depth into customer SCADA, and your regulatory-precedent corpus compound. The architectural reasoning sits in three senior heads, and a developer is offering two of them forty percent more. The moat sits underneath the model layer. The window is eighteen months.

03 · R&D CAPITAL ALLOCATION

Is our R&D budget one instrument or two?

Plateau capital runs the existing forecasting, asset-intelligence, trading, or diligence 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.

What the firms winning this play are actually doing

Four concrete moves any energy and environment 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 integration layer rotting**

Instrument review-time-per-line as a floor on any code touching customer SCADA, DCS, or CIM environments. PRs above five hundred lines require two senior reviewers and thirty minutes of documented review. Architectural decision records mandatory on agent-generated changes: no ADR, no merge. Evaluation harness coverage as a first-class metric on forecast models and asset-intelligence inference: golden set of balancing intervals, forecast-error rubric, automated grading, no harness, no ship.

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

Three programmes that stack over eighteen months. Asset-telemetry data-rights audit across your top customer contracts with a renegotiation plan tied to the next twenty-four months of renewals. A model-agnostic agent-memory store for asset behaviour, regulatory precedent, and integration knowledge, portable across foundation-model swaps. An integration-depth graph published for your top ten customer SCADA and DCS environments. Model layer commoditises. This layer compounds.

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

Do not migrate the whole book. Keep project-fee or capacity-fee for the installed base. Ship outcome-priced (per percent of forecast-error reduction, per megawatt-hour of avoided curtailment, per permit delivered on time) on a separate P&L for net-new and for customers explicitly asking. Write the cannibalisation trigger and kill criteria in Q1. 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 on controls, data science, and regulatory. Reverse it. Juniors in 2026 do not start on routine code. The agent layer absorbed that work. They start on integration audit, exception-handling, and paired composite work with a named senior for their first eighteen months. 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 energy and environment 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 and industrial-cloud professional services arms price the engagement as enablement for cloud or platform consumption. Their business model requires lock-in. Every build-or-buy decision routes through their pricing and their APIs for three years.

The AI-platform vendors and asset-intelligence upstarts have a quote on your desk. Their product is a feature-velocity intervention measured on seats or asset-connection count. Telling you to spend freed capacity on the data-rights and integration-depth moat instead is structurally impossible for them to say.

An in-house or fractional CAIO splits the seat. Strategy CAIO is revenue-measured; Platform CAIO is uptime-measured. One person cannot hold both in an asset-operations environment. 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. Asset-telemetry data-rights playbook. Regulatory-precedent query tool. 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.

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 starting to pitch 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 senior controls engineer or your lead regulatory specialist left for a developer at forty percent premium tomorrow, how many months would it take to reconstruct the integration and regulatory reasoning behind your core product?

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 your integration foundation 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 AI-strategy ownership list. Is your Q1 boundary agreement with the CEO written, or 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