

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 device sits above a model layer that commoditises every eighteen months. Your scorecard measures submission-cycle time and engineering headcount, and neither is the 2026 question. A briefing for CTOs, CPOs, and Heads of R&D at mid-market European medtech firms.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market medtech
technology and product leaders

DATE

April 2026

Monday 9:15, R&D review

Every CTO at a mid-sized European medtech firm has had the same Monday. Firmware sprint velocity up nineteen percent quarter on quarter. Submission-documentation agents cut GSPR assembly time by sixty percent. Your VP Engineering opens the review with the dashboard. He is mid-sentence when your phone buzzes. It is the CMIO at your biggest hospital system. One sentence: "We are piloting Aidoc on three imaging workflows, happy to talk about integration." Your Head of Regulatory Affairs mentioned a recruiter call from RegDesk over coffee Thursday. Your CEO has a CAIO shortlist from the board on his desk. Your CFO forwarded a Dexcom market-cap chart this morning: "Worth discussing before Q2 planning."

You have run this function for a decade or longer. The pattern used to be simple. Hire strong engineers. Clear the submission. Hold the notified body relationship. Ship the next device generation. Meet the board on cycle time and engineering headcount. The business was understandable.

Something changed. What used to take fifty engineers in firmware and submission prep 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 quality system actually absorbs is where your seat now lives or loses. An AI-native diagnostic at ten percent of your submission cost just landed its first named reference at a large hospital system in your category. Your Head of Regulatory Affairs carries notified body knowledge that has never been written down, and a regulatory-AI vendor 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 the next generation of the existing device through the existing regulatory pathway. The other funds what has to exist by 2028: the clinical-evidence engine you have not built, the installed-base telemetry platform you have not commercialised, the data rights you have not renegotiated with your top hospital systems, the architectural reasoning sitting in one RA director's head. 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.

Medtech just lost the equation that priced the category

The AI-natives in your category are not experimenting. They are clearing submissions. Paige.AI and PathAI shipped pathology platforms to 510(k) faster than your last device generation did. Aidoc and Rad AI cleared imaging-triage algorithms in months and priced them as subscriptions, not capex. RegDesk and Veeva Vault RIM AI are compressing what used to be a nine-month documentation exercise into weeks. Your hospital customers' procurement teams now run comparison quotes where your device capex number sits next to the AI-native's per-scan or per-patient subscription. They are not asking which is the better device. They are asking which pricing model their CFO can defend against the board's AI-readiness question.

The moat your product ran on for a decade has moved. The device-feature layer commoditises fast now. A diagnostic capability that used to take a 150-person team four years to build and submit can be rebuilt by a 20-person AI-native in twenty-four months, on modern architecture, priced on an outcome your hospital customer can measure. Your device-feature edge was never going to hold. What can hold is the layer underneath the sensor and the firmware: the installed-base telemetry, the clinical-data graph, the real-world evidence you have rights to aggregate across three product generations, the integration depth your service platform has wired into Epic and Cerner and DICOM workflow. Those compound. They take years to build and they cannot be bought. An AI-native cannot ship them because they do not have your installed base.

Your diagnostic cleared on the current AI tools is on an eighteen-month commoditisation curve. Your telemetry and clinical-data graph from ten thousand devices in the field is the asset the AI-native cannot replicate without your installed base.

The industry's speed mismatch is the third force. AI-native diagnostics build from scratch with no installed-base legacy, no three-year device-development cycle, no notified body relationship to protect. 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 portfolio, 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 deploying RegDesk 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. Firmware Copilot adoption is at seventy percent. Submission-assembly agents cut the GSPR documentation exercise by more than half. Your IEC 62304 evidence documentation lags the code. The review discipline that worked in 2022 is training the wrong practice on agent-generated volume now. Business-layer AI and regulated-submission AI are two different problems with two different controls, and the team is treating them as one.

Your platform and data. Whichever firm wins the next decade of medtech will own the data platform, not the sharpest sensor. Your installed base generates telemetry every day from thousands of devices in the field. Today that signal is thrown away or stored as service-contract ballast. Commercialise it. The compounding capital belongs in the clinical-data graph, the installed-base telemetry platform, the data-rights renegotiation across your top hospital contracts on the twenty-four-month renewal window, and the integration depth published for Epic, Cerner, and DICOM workflow. The device commoditises. The data platform compounds.

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 assumes AI-native diagnostics will price like device accessories. They do not. Paige and Aidoc price as subscriptions per patient or per scan. Your capex-priced P&L cannot follow. 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 junior engineers in 2024. The routine firmware and documentation work that trained juniors into seniors is now running on agents. No routine work, no training ground, no senior bench in 2030. Your Head of Regulatory Affairs is three years from retirement and the notified body relationship is not written down. 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-native diagnostics are shipping subscription-priced reference customers at ten percent of your submission cost. 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 medtech CTOs have not named yet.

The moat that matters in 2028 is not above the model layer. It is underneath it. The firmware feature you clear in this submission cycle on the current AI tools is on an eighteen-month commoditisation curve. The next foundation model rolls in 2027 and your diagnostic feature commoditises. What does not commoditise: the clinical-data graph your installed base generates, the telemetry signal from ten thousand devices in the field, the real-world evidence you have rights to aggregate, and the integration depth you have published for your top ten hospital IT systems. A two-person AI-native cannot build that layer because they do not have the installed 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 hospital footprint to start building the same layer.

From 2010 to 2020, three firms competed in continuous glucose monitoring. The first shipped the most integrated pump-plus-sensor system. The second shipped a half-price alternative with excellent industrial design. The third shipped a sensor that was clinically best-in-class on accuracy, and quietly built the data platform underneath it. By 2020 the third firm's market cap crossed forty billion dollars. The first firm's diabetes business was under strategic review. The 2010 scorecard measured sensor accuracy, sensor longevity, insertion comfort, and hardware reliability. Every one of those metrics stayed green at the first firm. They just stopped being the ones 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 submission-cycle time, and a scorecard that stopped mattering the quarter their biggest hospital account opened an AI-native 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

Deploy RegDesk, track GSPR savings

Regulatory documentation time cut sixty percent. IEC 62304 review lags the code. Business-layer AI and submission-layer AI blurred under one policy.

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 and submission cadence.

Ship AI features free into the existing device rate card

Hospital perceived value rises. Capex ASP flat. AI-native diagnostic lands your first reference loss at a subscription price your rate card cannot follow.

Cut the junior layer for Q1

Cheaper function now. Empty bench by 2031. External RA director at a premium.

THE WINNERS' PATH

Buy the tool layer, rebuild the review bar

Review-time floor per firmware change. Regulated-submission AI on validated pathway. Eval harness for evidence synthesis as first-class.

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 subscription-priced line on protected P&L

Parallel pricing with CFO and commercial. Capex holds installed base. Subscription wins AI-native net-new. No cannibalisation.

Pair juniors with seniors

Firmware juniors on architecture. RA juniors on notified body. Apprenticeship on judgment.

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 medtech 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 submission-grade output, or demos that fail the notified body?

Firmware Copilot at seventy percent. Submission-assembly sixty percent faster. Evidence documentation lags. A competitor reached 510(k) six months faster not because their team was smarter, but because their regulatory knowledge was systematic. Buy what the tool layer can carry. Own the review bar that keeps firmware and evidence submission-ready.

02 · PRODUCT DEFENSIBILITY

What does our portfolio do that an AI-native diagnostic at ten percent of our submission cost cannot copy?

Your device-feature edge commoditises every eighteen months when the next model rolls. Your installed-base telemetry and clinical-data graph compound. Your notified body relationship sits in one RA director's head, and a regulatory-AI vendor called last month. The moat sits underneath the device layer, not above it.

03 · R&D CAPITAL ALLOCATION

Is our R&D budget one instrument or two?

Plateau capital runs the existing device through the next submission. 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 authors the decade.

What the firms winning this play are actually doing

Four concrete moves any medtech 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 submission rotting**

Instrument review-time-per-firmware-change as a floor, not just cycle time. IEC 62304 classification check mandatory on every agent-generated change: no classification, no merge. Evaluation harness coverage for clinical-evidence synthesis as a first-class metric. Clear separation of business-layer AI (allowed) from regulated-submission AI (validated pathway). Rebalance the senior-to-junior ratio around the composite. The dashboard is no longer cycle time. It is review depth against submission-readiness.

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

Three concrete programmes that stack over eighteen months. Commercialise the installed-base telemetry pipeline as a data product separate from the device SKU. Build the clinical-data graph with data rights renegotiated across your top hospital contracts, tied to the next twenty-four months of renewals. Publish an integration depth graph for your top ten hospital IT systems — Epic, Cerner, DICOM workflow — maintained as structured metadata. The device commoditises. This layer compounds.

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

Do not migrate the whole portfolio. Keep capex for the installed base that does not want to switch. Ship subscription-priced diagnostic or AI-assisted service on a separate P&L for net-new hospital accounts and for customers explicitly asking. Write the cannibalisation trigger and kill criteria in Q1. The product architecture supports both mechanics natively. The joint Pricing Architecture Document with your CFO and commercial leader is signed before the first subscription-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 and the 2029 RA bench. Reverse it. Juniors in 2026 do not start on routine firmware or routine documentation. The agent layer absorbed that work. They start on architecture decisions, IEC 62304 compliance reasoning, and notified body relationship building, paired 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 medtech CTO has four places to go. None can do what we do.

The big health-practice consultancies ship a standardised AI-governance framework across a sector book. Their cost structure requires reusability. Your seat is idiosyncratic. A reusable framework on an idiosyncratic seat produces governance theatre.

The hyperscaler health-cloud arms price as enablement for cloud consumption. Their business model requires platform lock-in. Every build-or-buy on your telemetry and clinical-data graph routes through their pricing for three years.

The medtech regulatory-AI vendors have a quote on your desk. RegDesk, Veeva Vault RIM AI, Rimsys. Their product bills on submission volume. Telling you to spend freed capacity on the clinical-data moat instead is structurally impossible for them.

An in-house or fractional CAIO splits the seat. Strategy CAIO is revenue-measured; Platform CAIO is deployment-measured. Two-thirds of new CAIOs spend year one compiling AI inventories instead of setting direction.

WHAT YOU GET

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

Working infrastructure your team uses the same week

Custom agents built on your portfolio, 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 diagnostic 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 Regulatory Affairs, Head of Clinical Affairs, and Head of Service Platform each leave with five to ten pieces of working infrastructure tuned to their function. Submission-review audit agent. Notified body precedent tracker. Clinical-evidence synthesis harness. Telemetry platform 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 diagnostic or platform in your category. What is their pricing mechanic, and at what percent of your submission cost are they clearing? If you cannot name the top three, that is the first thing the briefing changes.

If your Head of Regulatory Affairs or two founding engineers left for an AI-native tomorrow, how many months would it take to reconstruct the notified body relationship and the architectural reasoning behind your core product decisions?

RegDesk or Veeva Vault RIM adoption at seventy percent. Submission-assembly time down sixty percent. Where did the freed hours go, and is your review bar still the one you want the firmware layer trained on?

Thirty percent CIO, twenty-three percent CDAO, twenty-one percent CAIO, one percent COO, zero percent CTO-CPO 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 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