



AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in the Moat You Build

Your R&D budget runs three programmes pretending to be one. Your application-engineering knowledge sits in three senior heads past sixty. Your scorecard measures release velocity and engineering headcount, and neither is the 2026 question. A briefing for CTOs, CPOs, and VPs of R&D at mid-market European B2B industrial firms.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market B2B
industrial
technology and product leaders

DATE

April 2026

Monday 9:15, R&D review

Every CTO at a mid-sized European industrial firm has had the same Monday. Q1 hardware release shipped two sprints early. Application engineering up twenty-eight percent on billable hours. Service platform ingesting eleven million telemetry events a day. Your VP Engineering opens the review with the dashboard. He is mid-sentence when your phone buzzes. It is your largest distributor. One line: "A customer ran your rate card through Coupa against three European competitors and a Tulip-on-open-hardware build. We're losing on price and on lead time. Thursday?" Your top application engineer sent a calendar invite titled retirement planning. Your CEO has a Chief Digital Officer shortlist from the board on his desk. Your CFO forwarded the Siemens Digital Industries earnings this morning. Worth discussing before Q2 planning.

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

Something changed. What used to take fifty engineers on a product-development team now runs on five, and the same compression lands whether your team ships a pressure sensor, a bearing, a valve, or a complete automation cell. Only the denominator changes. 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. A software-native on open hardware just landed its first named reference customer in your category. Your senior application engineers carry the reasoning behind every successful install of the last decade, and two of them are past sixty.

You are not running one R&D function. You are running three, and only one of them is on your scorecard. One funds the next hardware release and the catalogue extensions. One funds the commercial-ops tech stack your sales and application-engineering teams live on. One funds the service and telemetry platform. All three pull from the same senior engineers. Only one of the three builds the 2028 moat: the application-engineering knowledge captured, the installed-base telemetry structured, the integration depth published to your top ten customer systems. On one scorecard the other two do 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 industrial just lost the equation that priced the category

The software-natives in your category are not experimenting. They are shipping. Tulip, Samsara, Cognite each own a slice of what you thought was an application-engineering moat, and each is pricing on outcomes — uptime hours, units produced, defects avoided — at a fraction of your per-unit plus service-contract economics. Your enterprise customers' procurement teams run comparison quotes where your number sits next to the software-native's outcome number and, increasingly, next to an open-hardware build the customer can spec through a Coupa or LevaData benchmarking tool in eleven minutes. 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 component spec layer commoditises fast now. An AI-armed procurement buyer does not need a distributor to tell them what sensor fits the line, or what valve fits the flow. What can hold is the layer underneath the catalogue: the application-engineering knowledge captured from the last twenty years of customer installations, the telemetry from the installed base your service contracts do not yet give you rights to, the integration depth you have wired into the top customer PLC and DCS systems. Those compound. They take years to build and cannot be bought. A software-native on open hardware cannot ship them because they do not have your installed base.

Your component shipped on the current catalogue is on an eighteen-month commodity curve. Your integration depth with your top ten customer systems is the asset the software-native cannot replicate without your installed base.

There is a second force in the industry that most mid-market CTOs have seen play once already. GE bet Predix. Siemens built MindSphere. Both industrial giants spent a decade convinced the moat was the platform layer on top of the installed base. Both platforms worked technically. Neither captured the moat. The moat was always the application-engineering knowledge already in the engineers' heads plus the telemetry nobody had structured yet. A decade later the moat is still there, waiting, for the firm that builds underneath the model layer instead of on top of it. The firms that are already moving are quiet about it. The ones that are piloting another predictive-maintenance demo 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. Your hardware team is running design-space exploration in days that used to take months. Firmware is shipping on agent-generated test suites. Your application engineers are drafting commissioning plans from customer plant-floor photos. Velocity is up double-digit on the dashboard. And somewhere in the same quarter a field-failure root cause turned out to be a simulation corner the agent did not test and nobody senior reviewed. The review discipline that worked in 2022 is training the wrong practice on agent-generated volume now.

Your platform and data. Your commercial-ops stack, your product-engineering platform, and your service-and-telemetry platform are three programmes funded separately, staffed from the same senior pool, and integrated by nobody. Every industrial giant that spent a decade on an IIoT platform bet learned the moat was never the platform. It was the application-engineering knowledge and the installed-base telemetry underneath. You cannot replicate a decade of platform investment 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: an application-engineering knowledge graph extracted from work artifacts, installed-base telemetry with data rights renegotiated into service contracts, an integration depth graph published for your top ten customer systems.

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 in the product will price like legacy features. They do not. Tulip and Samsara are shipping outcome-priced on uptime. Your competitor in the next European country over just announced a per-defect-avoided contract. You are pricing on per-unit and per-service-hour. 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 software-native forces the conversation on their terms.

Your R&D bench. You stopped hiring juniors in 2024. The routine application-engineering work that trained juniors into senior application engineers is now running on agents. No routine work, no training ground, no senior bench in 2030, no internal CPTO candidate the year after. Your three senior application engineers carry the reasoning behind every successful install of the last decade. Two of them are past sixty. The retirement calendar invite from your top engineer is the signal the firm cannot afford to miss.

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. In industrial, where the product-engineering seat and the commercial-ops seat sometimes sit with different people, the authority is drifting twice as fast. Your CEO has a Chief Digital Officer 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 and COO lands before the CDO shortlist, not after.

The moat that matters in 2028 is not above the model layer

Your sector is repricing from below because software-natives on open hardware are shipping outcome-priced products and AI-armed procurement teams are commoditising the catalogue layer. Your function is compressing because five engineers now produce what fifty did, and the design-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 industrial CTOs have not named yet.

The moat that matters in 2028 is not above the model layer. It is underneath it. The next product release you ship this year on any foundation model is on an eighteen-month commodity curve. What does not commoditise: the application-engineering knowledge you have captured, the installed-base telemetry you have structured, the integration depth you have published for your top ten customer PLC and DCS systems. A software-native on open hardware cannot build that layer because they do not have the installed base that generates the signal. You do. An AI-armed procurement team cannot commoditise the moat because a rate-card comparison cannot price application expertise that lives in structured infrastructure your competitor does not have. The window closes in roughly eighteen months, which is the time between now and when a software-native has enough installed-base footprint to start building the same layer.

In 2015, the CEO of a century-old industrial giant bet the company on an industrial Internet of Things platform. Every metric inside the bet looked right. Industrial data growing exponentially. Predictive maintenance demonstrably valuable. Customer demand for analytics. The largest installed base in the world. A European competitor launched the same thesis a year later with a different architecture. By 2020, both platforms had been quietly shelved or spun off. Neither captured the moat. The moat was always the application-engineering knowledge already in the heads of the senior engineers, plus the telemetry from the installed base that nobody had structured yet. Both firms built on top of the knowledge instead of capturing it. A decade later the moat is still there, waiting.

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 one programme, or three programmes pretending to be one when they are three? 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 Chief Digital Officer hire arrives inherit a seat that has narrowed to platform uptime and release velocity, and a scorecard that stopped mattering the quarter their largest distributor lost a customer to a Tulip build.

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

Every industrial 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 tools, track adoption

Simulation and CAD-assist rolled out. Velocity up. Field-failure surprises compound. Review bar trains on agent volume.

Pilot another predictive-maintenance app

Third industrial-IoT demo this year. Board deck reads well. Moat is no closer. Senior engineers still unstructured.

Hold per-unit pricing, add aftermarket services

Rate card anchored in components. Customer value rises thirty percent. ARPU flat. Tulip lands the next reference loss.

Cut the junior layer for the Q1 headcount win

Cheaper function this quarter. Senior bench retires in 2027. External hire at a premium the firm cannot absorb.

THE WINNERS' PATH

Buy the tool layer, redesign the review bar

Review-time floor per firmware module. Design-review mandatory on agent-generated changes. Eval harness as infrastructure.

Capture knowledge, structure telemetry, publish depth

Application-engineering graph from existing artifacts. Installed-base data rights renegotiated. Integration depth to top ten customer systems.

Stand one outcome-priced line on protected P&L

Per-unit holds the installed base. Outcome-priced wins net-new. Joint Pricing Architecture Document with CFO and commercial leader.

Capture knowledge before retirement, pair juniors

Top five senior engineers' reasoning encoded in twelve months. Juniors pair on real applications. 2031 bench rebuilt deliberately.

Same engineering bench. 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 industrial 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 at commissioning?

Hardware cycles compressed by simulation. Firmware on agent-generated test suites. Application engineering on customer-documentation agents. Velocity is up. Field-failure root causes are up too. The industrial giants spent a decade building the platform foundation that lets application teams ship at speed. You cannot replicate a decade 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 Tulip on open hardware or an AI-armed buyer cannot commoditise?

Your component spec commoditises when the next AI-armed procurement tool runs a benchmark. Your application-engineering knowledge, your installed-base telemetry, and your integration depth compound. Your top senior engineers carry the reasoning behind every successful install, and two of them are past sixty. The moat sits underneath the model layer, not above it. The window to capture it is eighteen months.

03 · R&D CAPITAL ALLOCATION

Is our R&D budget one programme or three?

Hardware R&D funds the next release and catalogue extensions. Commercial-ops R&D funds the sales and application-engineering tech stack. Service and IoT R&D funds the telemetry platform. On one hurdle rate, hardware wins every quarter. On one scorecard, the other two do 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 three, each defended separately, authors the decade.

What the firms winning this play are actually doing

Four concrete moves any industrial 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 field-failure surprises**

Instrument review-time-per-module as a floor, not just cycle time. Hardware changes above a complexity threshold require two senior reviewers and a documented design-review. Firmware agent-generated test suites run against a golden-set regression before merge: no regression pass, no ship. Commissioning checklists are agent-drafted and senior-reviewed, never agent-signed. The dashboard is no longer release velocity. It is review depth against agent-generated output volume.

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

Three programmes that stack over eighteen months. An application-engineering knowledge graph extracted from work artifacts your seniors already produce: customer-site reports, commissioning records, call transcripts, warranty-claim histories. No SharePoint graveyard. Installed-base telemetry structured with data rights renegotiated across your next twenty-four months of service-contract renewals. An integration depth graph published for your top ten customer PLC, DCS, and SCADA systems. Model-agnostic across swaps.

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

Do not migrate the whole book. Keep per-unit and per-service-hour for the installed base that does not want to switch. Ship outcome-priced on uptime, units-produced, or defects-avoided for net-new and for service-contract renewals already asking. Write the cannibalisation trigger in Q1: if outcome revenue on an account exceeds per-unit revenue, trigger the migration conversation. The joint Pricing Architecture Document with your CFO and commercial leader is signed before the first outcome-priced deal closes.

4**R&D bench: capture senior knowledge before retirement, pair juniors deliberately**

Your top five senior application engineers are past fifty-five. Two are past sixty. Their reasoning is the firm's moat and it is a retirement decision away from walking out. In the next twelve months, capture it from existing work artifacts. The senior reviews the agent-extracted graph and corrects it. Juniors in 2026 pair with a named senior on real applications for eighteen months. Your 2031 bench is rebuilt deliberately instead of hollowed out by default.

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

The industrial 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 IIoT platform vendors price the engagement as enablement for platform consumption. Their business model requires platform lock-in. Every build-or-buy decision routes through their pricing for three years. The Predix and MindSphere decade is the cautionary tale.

The industrial automation and digital-manufacturing vendors have a quote on your desk. Rockwell + PTC, Siemens Digital Industries, Tulip, Samsara, Cognite. Their product is a platform intervention measured on platform penetration. Telling you to capture the moat underneath their platform instead of inside it is structurally impossible for them to say.

An in-house or fractional Chief Digital Officer or CAIO splits the seat. Commercial-side CDO and industrial-IIoT Head carry incompatible mandates. One person cannot hold both. Two people create parallel governance and duplicate the R&D capital conversation.

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 products, your customers, your stack

The start of a knowledge-and-telemetry 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 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 across hardware, commercial-ops, and service and IoT, and runs the alternatives analysis against your named software-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 Application Engineering, and Head of Service and IoT each leave with five to ten pieces of working infrastructure tuned to their function. Application-engineering knowledge-graph extractor. Design-review audit agent. 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 three programmes with three hurdle rates and three kill criteria. The boundary agreement with the CEO, CFO, and COO 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 programme or three? If one, what does the board see that splits hardware, commercial-ops, and service and IoT, and what is the kill criterion on each?

Name the software-native in your category on open hardware. 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 three top senior application engineers retired in the next eighteen months, how many months would it take to reconstruct the application-engineering reasoning behind your last decade of successful installs?

Your service contracts run for multiple years. What portion of your installed-base telemetry do you have data rights to, structured, and usable? What gets renegotiated in the next twenty-four months of renewals?

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, CFO, and COO written, or waiting until the Chief Digital Officer 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