

What's Changing Under Your Commercial Platform

Your commercial-platform roadmap is twenty-four months deep. A Gong quote is on your desk. Your top technical seller retires in January. Your scorecard measures uptime and release cadence, and neither is the 2026 question. A briefing for CTOs, CPOs, and Heads of Digital at European industrial manufacturers with sales-heavy commercial organisations.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European industrial
manufacturers
with technical sales organisations

DATE

April 2026

Monday 9:15, Q2 platform review

Every CTO at a sales-heavy industrial manufacturer has had the same Monday. Salesforce upgrade on track. CPQ release cadence steady. Field-service mobile rollout at eighty percent of sites. Your VP Platform Engineering opens the review with the dashboard. Your phone buzzes mid-sentence. Email from your VP Sales: "Gong demo Thursday. CFO wants to join." Ten minutes later, another from your Head of Applications: "Bengt retires in January. Nobody on the bench carries his petrochemical knowledge." Your CEO forwarded a Pepperl+Fuchs case study this morning with one line: "Worth discussing before Q3 planning."

You have run this function for ten years. The pattern used to be simple. Keep the CPQ, CRM, and field-service platforms running. Ship the roadmap. Hire strong platform engineers. Meet the board on uptime and release cadence. The business was understandable.

Something changed. What used to take fifty engineers to ship a CPQ release 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. Your 1,000 field sellers spend forty to sixty percent of their time preparing to sell and documenting what was sold. Your top three technical sellers retire between 2027 and 2029, and their application knowledge lives in their heads, not in your CPQ.

You are not running one commercial-platform function. You are running two, and only one of them is on your scorecard. One keeps the CPQ, CRM, and field-service platforms running cheaper. The other builds what has to exist by 2028: the seller-intelligence layer underneath the platforms you already pay for, trained on your installed base, your configurator rules, your call transcripts. 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 platform organisation you run, and the three questions your team should answer before the next board cycle.

Industrial sales just lost the equation that priced the catalogue

Two forces are compressing your commercial organisation from opposite sides. On the buyer side, procurement teams are arming themselves with AI. LevaData, Coupa, and Keelvar are running part-level cost benchmarking and supplier evaluation at fifty-eight percent production deployment in industrial procurement. A seller who walks into a quote conversation without the installed-base depth, the application knowledge, and the competitive context the buyer has already run through an AI platform does not get a second meeting. The buyer already knows your competitor's price. The buyer already has a value analysis. Your seller is closing the gap in forty minutes of in-meeting time with a product catalogue and good intentions.

On the seller side, a Gong quote is on your desk. A Seismic quote is on your desk. A Highspot quote is on your desk. Your VP Sales wants all three. Each is built for generic B2B SaaS sales cycles, not for a twenty-five-thousand-SKU catalogue with configurator rules, application constraints, and twelve-country pricing architectures. Each charges three hundred Euros per seat per month. For a 1,000-seller organisation that is €3.6M annually flowing to vendors whose product takes your call transcripts, your CRM data, and your installed-base records as training material. Over five years, €18M flowing out of an already-thin operating margin, and the seller-intelligence layer you wanted is still not yours.

The commercial-SaaS bundle you sign this quarter is the training data you hand your category's AI-procurement counterparty two years from now.

The firms ahead of this are quiet about it. They are not adding vendors to the stack. They are structuring their installed-base data, extracting application knowledge from call transcripts, tying it to configurator rules, and building their own seller-intelligence layer on top of the commercial platforms they already own. A 1,000-seller organisation at Pepperl+Fuchs ran the math on fifteen pilot sellers: €150K revenue responsibility each, a five-percent improvement, thirty-two times ROI on the programme. That math scales with every seller who touches the platform. The math on Gong across a thousand sellers at three hundred Euros per seat per month scales in the other direction.

Your platform function is feeling it every day

Your build engine. Cursor adoption across your platform team is at seventy percent. CPQ release cadence is steady but not faster. Agent-generated configurator rules are landing with reviews barely deeper than a rubber stamp. The foundational layer rots quietly. Twelve months in, your team cannot execute a platform rewrite when one arrives. The review discipline that worked on hand-written CPQ logic is training the wrong practice on agent-generated volume now.

Your platform and data. Your installed-base data sits in four systems, none of them structured for the way sellers actually ask questions. Your call transcripts flow through Gong into a training set you do not own. Your configurator rules encode thirty years of application learning in a format only three people understand. Spotify spent four years on the platform foundation that lets agents ship cleanly to production. 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: structured installed-base data, a call-transcript extraction pipeline tied to configurator rules, an application-knowledge graph published for your top product families. The model layer commoditises. The data layer compounds.

Your product thesis. Your VP Sales wants Gong. Your CFO wants the unit cost of the commercial platform down. Your CEO wants the Pepperl+Fuchs thirty-two-times ROI conversation this quarter. Three different asks. One stack. Your rate card assumes AI features price like legacy features. They do not. The build-vs-buy call on the seller stack is a capital-allocation question, not a procurement one. Your VP Sales and your CFO need a joint document with you before the commercial-SaaS vendors force the conversation on their terms.

Your R&D bench. You stopped hiring junior application engineers in 2024 because the routine tender-response work ran on agents. No routine work, no training ground, no senior application engineer bench in 2030 when Bengt and his two peers retire. OpenAI and Anthropic 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 application-engineering bench the year Bengt's last petrochemical customer calls.

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. Every month without a visible claim, the seller stack drifts further from the seat that can actually defend it. Your CEO has a CAIO shortlist. Your VP Sales has a Gong quote. Your Q1 boundary agreement with the CEO on who owns the seller stack lands before the CAIO shortlist, not after.

The platform that matters in 2028 is not the CRM layer above it

Your sector is being squeezed from both sides. AI-armed procurement compresses the buyer conversation before your seller walks in. Commercial-SaaS bundles compress the seller stack by claiming the intelligence layer. Your function is absorbing its own compression at the same time. Five platform engineers now ship what fifty used to ship, and the review bar built for fifty cannot hold the volume five produce with agents. Same force, two altitudes. But there is a third thing happening that most CTOs at industrial manufacturers have not named yet.

The platform that matters in 2028 is not the CRM layer above it. It is the seller-intelligence layer underneath it. A Gong plus Seismic plus Highspot bundle ships on the layer above. The next foundation-model release commoditises that layer. What does not commoditise is the application knowledge your top technical sellers carry, the configurator rules you have accumulated over thirty years, the installed-base data that encodes every sensor failure mode you have ever fixed, the call-transcript corpus you own. A commercial-SaaS vendor cannot build that layer because they do not have your product catalogue, your customer base, or your twenty-five thousand SKUs. You do. The window closes in roughly eighteen months.

In 2015, pharma companies had three options for CRM. Retrofit generic Salesforce with consulting and point tools. Build their own pharma-specific layer. Adopt Veeva. The firms that retrofit paid for licences plus consulting plus workflow gaps for a decade. The firms that built their own got the compounding advantage. The firms that adopted Veeva early got the platform moat without the build cost. By 2020 generic CRM was no longer a serious option for pharma field medical. The industry-specific layer had won. That decision was made in 2015, not in 2020. By 2020 the market had already moved.

You can tell which camp a firm is in from a single question. Is the commercial-platform 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 industrial manufacturers cannot answer. The ones that can have already made the shift.

You have twelve to eighteen months. The CTOs who claim authorship of the seller stack in Q1 author the decade. The ones who wait inherit a commercial-SaaS bill of €3.6M annually and a scorecard that stopped mattering the quarter their biggest competitor started extracting their shared call transcripts out of a Gong instance.

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

Sign the commercial-SaaS bundle

Gong plus Seismic plus Highspot at €300/seat/month across 1,000 sellers. Your call transcripts feed someone else's training set.

Route the seller stack through VP Sales procurement

Vendor selection by seat penetration and demo quality. CTO seat narrows to platform uptime. CAIO arrives in Q3.

Ship Cursor across the platform team

Adoption at seventy percent. CPQ release cadence flat. Agent-generated configurator rules landing with shallow review.

Cut the junior application-engineering layer

Cheaper function this quarter. Empty bench when Bengt retires in 2027. External hire at a premium you cannot absorb.

THE WINNERS' PATH

Build the seller-intelligence layer underneath

Structured installed-base data. Call-transcript extraction tied to configurator rules. Application-knowledge graph you own.

Route the build-vs-buy call through the capital-allocation hurdle

Joint document with CFO and VP Sales. Two budgets, two hurdle rates, two kill criteria.

Absorb the compression, redesign the review bar

Review-time floor on agent-generated rules. Mandatory application-engineer sign-off. Eval harness on configurator coverage.

Pair juniors with seniors on compounding work

Eighteen-month apprenticeship on application reasoning, not three years on routine tender drafts.

Same platform team. Same twelve months. Opposite outcomes. The difference is whether the CTO authors the seller stack or signs the requisition.

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

Three questions every industrial-manufacturer CTO is tracking. The second is the crux. The other two are how you earn the right to answer it.

01 · BUILD VELOCITY

Is our platform team shipping production-grade work, or volume that fails in the field?

Cursor adoption at seventy percent. CPQ release cadence steady. Agent-generated configurator rules 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 to production.

02 · PRODUCT DEFENSIBILITY

What does our seller stack own that Gong and Seismic cannot copy?

Your feature edge on generic tools commoditises every foundation-model release. Your installed-base data, your configurator rules, and your call-transcript corpus compound. Your application knowledge sits in three senior engineers' heads, and two retire before 2029. The moat sits underneath the vendor layer. The window to build it is eighteen months.

03 · R&D CAPITAL ALLOCATION

Is our commercial-platform budget one instrument or two?

Plateau capital keeps the existing stack running cheaper. Compounding capital builds the 2028 seller-intelligence moat. On one hurdle rate the first wins every quarter. On one scorecard the second does not exist. The CTO who walks in 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 industrial-manufacturer 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 agent-generated configurator rule as a floor, not just release cadence. Any agent-drafted tender response above a complexity threshold requires named application-engineer sign-off. Evaluation harness coverage on configurator changes and installed-base extractions as a first-class metric. Rebalance the senior-to-junior platform-team ratio around the composite. The dashboard is no longer release cadence. It is review depth against output volume.

2**Platform and data: build the seller-intelligence moat underneath the vendor layer**

Three programmes that stack over eighteen months. Installed-base data audit across your top product families, published as structured metadata your sellers can walk into procurement with. A call-transcript extraction pipeline tied to configurator rules, so application knowledge flows from the seller's customer visit back into your catalogue. An application-knowledge graph the firm owns. The vendor layer commoditises. This layer compounds.

3**Product thesis: joint build-vs-buy document with VP Sales and CFO**

Do not sign a commercial-SaaS bundle without a written data-rights and cannibalisation analysis. Cursor across your platform team is a productivity decision. Gong across your sales force is a moat-ownership decision. They are not the same conversation. Keep generic enablement where it genuinely earns its seat. Build the seller-intelligence layer underneath and own it. Joint Pricing Architecture Document signed before the first bundle renewal.

4**R&D bench: rebuild the application-engineering pathway before the top three retire**

Your 2024 decision to stop hiring junior application engineers saves Q1. It empties the 2030 senior bench. Reverse it. Juniors in 2026 start on audit, exception-handling, and paired application-reasoning work with a named senior for eighteen months, not on routine tender responses. 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 because it did not exist anywhere else.

The only firm that joins platform authorship to working infrastructure

The industrial-manufacturer 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 commercial-SaaS bundles have quotes on your VP Sales' desk. Gong, Seismic, Highspot, Outreach, Clari. Their product is a seat-billed layer above the CRM. Telling you to build the seller-intelligence layer underneath their own product 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 build-vs-buy conversation the seller stack 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 CPQ, your configurator, your installed base

The start of a seller-intelligence layer 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 platform leadership. Working infrastructure from week one.

We start with your platform and commercial 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 platform function. Build the first piece of infrastructure in the room: a seller-intelligence agent that takes a real customer file, pulls installed-base records, and drafts the application-specific value conversation. Your leadership team leaves with a shared view and the first working artefact already in hand.

WORKSHOP 02

Your seller stack

A half-day build. Your VP Platform Engineering, Head of Applications, Head of Data, and VP Sales each leave with five to ten pieces of working infrastructure. Pre-meeting intelligence agent. Configurator-rule audit scaffolding. Call-transcript extraction prototype. Tender-response acceleration agent. 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 commercial-platform budget into two instruments with two hurdle rates and two kill criteria. The boundary agreement with the CEO on who owns the seller stack. The joint Pricing Architecture Document with your VP Sales and CFO. The eighteen-month seller-intelligence build plan.

Three weeks. No IT project. No board approval needed to start. The infrastructure your commercial-platform team is already running on when we finish.

For your next platform review

The questions where your platform and commercial leadership cannot agree on an answer are the ones worth an hour on the agenda.

Is your commercial-platform 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-procurement counterparty at your biggest customer. What are they running, what part-level benchmarks are they feeding into their value analysis, and what does your seller walk in with that the buyer's system does not already have?

If your top three technical sellers retired between 2027 and 2029, how many months would it take to reconstruct their application knowledge from what your commercial platform captures today?

Cursor adoption on your platform team is at seventy percent. CPQ release cadence is flat. Where did the freed hours go, and is your review bar still the one you want the configurator 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 on who owns the seller stack written, or is it waiting until after the CAIO shortlist lands?

If any of these produces a disagreement in your platform and commercial 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