



AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in the Stack You Own

Your commercial-AI budget is two instruments pretending to be one. Your HCP intelligence sits underneath a vendor layer that commoditises every upgrade cycle. Your scorecard measures rep adoption and integration backlog, and neither is the 2026 question. A briefing for commercial CTOs, Heads of Commercial Technology, and VPs Commercial Platforms at mid-sized European pharma affiliates.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-sized pharma
commercial technology leaders

DATE

April 2026

Tuesday 10:40, affiliate digital steering

Every commercial CTO at a mid-sized European pharma affiliate has had the same Tuesday. Rep adoption of Veeva's Pre-call Agent at seventy-three percent. MLR cycle time down to thirty-six hours on standard content, still forty-eight on Medical Affairs work. Your Head of Commercial Ops is mid-slide when your phone buzzes. It is an email from your Regional Digital Director: "Global is bringing the CAIO hire in front of the executive committee Thursday. Standardisation recommendation on commercial-AI architecture across affiliates. Wanted you to have a heads-up." Your Veeva account manager has the Q3 AI module expansion proposal on your desk. Your Salesforce Life Sciences migration plan is eighteen months into a twenty-four-month runway. Your head of MDM mentioned a recruiter call from Accenture Life Sciences over coffee Thursday.

You have run this stack for ten years. The pattern used to be simple. Keep Veeva stable. Integrate IQVIA cleanly. Ship the roadmap. Meet the affiliate GMT on uptime and rep adoption. The business was understandable.

Something changed. Veeva shipped Pre-call Agent, Voice Agent, Free Text Agent in December 2025, with April, August, and December 2026 waves coming. Salesforce Life Sciences Cloud went generally available in October 2025. IQVIA OCE+ embedded Next Best Action in production. What used to be your integration roadmap is now a vendor-AI subscription ladder landing on the SG&A page, not a pilot you can stage. Ninety-five percent of corporate AI deployments produce no measurable P&L impact. The gap between what the tools ship and what your affiliate's system absorbs is where your seat now lives or loses.

You are not running one commercial-AI stack. You are running two, and only one of them is on your scorecard. One funds what Veeva, Salesforce, and IQVIA sell back cheaper through more AI copilot seats. The other funds what has to exist by 2028: the HCP interaction corpus you have not captured, the consent-scope graph you have not drawn, the agent-memory store that persists MSL context across vendor upgrades, the integration depth with Vault, PromoMats, and pharmacovigilance intake that outlasts three model swaps. On one licence-cost hurdle 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 commercial-tech function you run, and the three questions your team should answer before Q3 locks.

Pharma commercialisation just lost the equation that priced the category

The vendor AI modules in your category are not experimenting. They are shipping. Veeva rolled Pre-call Agent, Voice Agent, Free Text Agent in December 2025, with April, August, and December 2026 waves on the roadmap. Salesforce Life Sciences Cloud went generally available in October 2025 with seventy-plus global pharma customers signed on the wave. IQVIA OCE+ shipped embedded Next Best Action in production. The AI copilot inside your CRM and your next-best-action engine is now a line on every affiliate's SG&A page in every market you compete in, on the same upgrade cycle, grounded on the same cross-affiliate aggregated layer.

The moat your commercial tech ran on for a decade has moved. The productivity layer commoditises fast now. A capability that used to take your team three years to build into the stack (Next Best Action routing, Pre-call summaries, Content first-pass review) ships to every affiliate in the category by the next Veeva or Salesforce release, on a modern architecture, priced on the licence line your GM already signs. Your productivity edge was never going to hold. What can hold is the layer underneath the vendor stack: the HCP interaction corpus your affiliate has been generating for ten years, the consent-scope records your Medical Affairs and Market Access teams own, the MSL-to-KOL interaction chain that lives in twelve systems and has never been captured as institutional infrastructure. Those compound. They take years to build and they cannot be bought. A vendor cannot ship them because they do not have your affiliate's HCP reality.

Your Pre-call Agent summary is on an eighteen-month obsolescence curve. Your HCP interaction corpus and consent-scope graph are the assets the vendor AI module cannot replicate without your affiliate's ten-year data trail.

The industry's speed mismatch is the third force. Veeva, Salesforce, and IQVIA ship AI on their release cycle, not yours. Global HQ is about to push a standardised commercial-AI playbook down through the CAIO hire. EU Joint Clinical Assessments since January 2025 are demanding real-time cross-market evidence coordination that affiliates were not architected for. Your affiliate has all three pressures. You cannot pivot at vendor speed on vendor terms. What you can do is identify what is genuinely defensible in your commercial-tech stack, stop defending what is not, and ship the institutional infrastructure the vendor AI cannot reach. The affiliates that are already moving are quiet about it. The ones that are buying Q3 AI add-ons and calling it a commercial-AI strategy are the ones the Regional VP will rebid the platform from in 2028.

Your commercial tech is feeling it every day

Your build engine. Pre-call Agent adoption is at seventy percent across the field force. MLR cycle time has barely moved on Medical Affairs content. Content Agent compresses first-pass review, not the Medical-Director judgment underneath it. Your integration backlog is forty tickets deep on the Salesforce Life Sciences migration. Velocity is green on the monthly pack and the pharmacovigilance validation team is slipping on agent-generated intake. The discipline that worked in 2022 is training the wrong practice under vendor-AI volume now.

Your platform and data. Spotify's best developers have not written code since December because Spotify spent four years on platform foundation. Your affiliate did not spend four years. Global HQ did not fund it. You can buy the build-side discipline from Veeva, Salesforce, and IQVIA and redirect the compounding capital into the layer the vendors cannot sell: the HCP interaction corpus captured as institutional infrastructure, a consent-scope graph drawn across channels, a model-agnostic agent-memory store portable across Veeva and Salesforce and IQVIA upgrade cycles, an integration depth graph published for Vault, PromoMats, pharmacovigilance intake, and transfers-of-value filing. The vendor layer commoditises. This layer compounds.

Your product thesis. Global asked for the affiliate's three-sentence commercial-AI thesis last quarter. Your answer was the one from eighteen months ago. The CAIO deck the Regional VP sees on Thursday assumes your affiliate's job is to deploy the vendor playbook cleanly. That is the authorship question, not the deployment one. Your Medical Director, your Compliance Officer, and your GM need a joint document with you naming which layers the affiliate owns and which are vendor substrate, before the CAIO deck decides for you.

Your platform bench. You stopped backfilling junior commercial-tech ops in 2024. The routine monitoring and integration work that trained juniors into senior MDM stewards is now running on agents. No routine work, no training ground, no senior bench in 2030, no internal architect-seat candidate the year after. The pattern that OpenAI and Anthropic restarted in late 2025 (pairing juniors with seniors on judgment-intensive work from day one) applies here too. The affiliate that cuts the junior layer for this year's SG&A 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. Commercial CTOs are not on the list. Not because commercial-tech leaders abandoned the mandate. Most have not yet claimed it. Your Global CEO has a CAIO hire and a standardisation deck. Every month without a visible claim, the thesis drifts further from the seat that can defend it. Your Q1 boundary agreement with the GM lands before the CAIO deck, not after.

The moat that matters in 2028 is not inside the vendor AI module

Your sector is repricing from below because the vendor AI modules ship the same capability to every affiliate in the category on the same upgrade cycle. Your function is compressing because rep adoption of Pre-call Agent looks green while MLR cycle time on Medical Affairs content will not move. Same force, two altitudes. But there is a third thing happening that most mid-sized affiliate commercial CTOs have not named yet.

The moat that matters in 2028 is not inside the vendor AI module. It is underneath it. The Pre-call Agent is grounded on something. The Next Best Action runs on something. The Content Agent reviews against something. That something is either your affiliate's institutional infrastructure that you own, or the vendor's cross-affiliate aggregated layer that they own. For most affiliates today, the answer is the vendor's. The compounding move is to rebuild the grounding layer as affiliate-owned infrastructure while the vendor still accepts affiliate-grounded inputs. A two-person AI-native cannot build it either. They do not have ten years of your HCP reality. You do. The window closes in roughly eighteen months.

In 1995, a specialty medical publisher shipped the best journals in its category. The subscription line was healthy. The Editor-in-Chief's scorecard was article volume, citation index, peer-review rigour. Every metric was green. A new kind of product shipped around the decade turn: UpToDate, DynaMed, ClinicalKey. Not better journals. A platform that aggregated what the category had always produced and had never captured as institutional infrastructure: the citation graph, workflow integration, the reference layer inside clinical settings. By 2015 the revenue migration was visible. The 1995 scorecard never turned red. It stopped being the one the market was scoring.

You can tell which camp an affiliate is in from a single question. Is the commercial-AI budget you signed this quarter a plateau budget, a compounding budget, or one budget pretending to be one when it is two?

You have twelve to eighteen months. The commercial CTOs who claim authorship of this decision in Q1 author the decade. The ones who wait until the CAIO deck lands inherit a seat that has narrowed to uptime and integration backlog.

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

Every commercial CTO has had the vendor-AI conversation. Most are making the same four decisions. A smaller group is making different ones.

THE DEFAULT PATH

Sign the Q3 Veeva AI expansion, track rep adoption

Seventy-percent adoption. MLR queue flat on Medical Affairs content. Pharmacovigilance validation slipping on agent-generated intake.

Route commercial-AI spend through the vendor licence line

Global CAIO owns the architecture decision within eighteen months. Affiliate seat narrows to uptime and integration backlog.

Let vendor AI ground on cross-affiliate aggregated data

Pre-call summaries and Next Best Actions pulled from vendor layer. Your HCP reality leaks upward, returns commoditised.

Cut the junior commercial-tech bench for this year's SG&A

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

THE WINNERS' PATH

Buy the vendor-AI layer. Redesign the validation bar.

Mandatory review gates on agent-generated Medical Affairs output. Eval harness on Pre-call Agent grounding against affiliate HCP reality.

Route commercial-AI capital through a split hurdle

Two budgets, two hurdles, two kill criteria. Plateau runs the vendor stack cheaper. Compounding builds the 2028 HCP data moat.

Ground vendor AI on affiliate-owned corpus

Published HCP interaction layer. Consent-scope graph drawn. Affiliate grounding accepted into the vendor AI stack while the window is open.

Pair juniors with senior MDM and integration leads

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

Same commercial-tech team. Same twelve months. Opposite outcomes. The difference is whether the CTO authors the commercial-AI capital allocation or signs the vendor requisition.

Build Velocity. Product Defensibility. Commercial-AI Capital Allocation.

Three questions every commercial CTO at a mid-sized European pharma affiliate is tracking. The third is the crux. The first two are how you earn the right to answer it.

01 · BUILD VELOCITY

Is our commercial-tech speed shipping validated work, or demos the Medical Director rejects?

Pre-call Agent adoption at seventy percent. MLR cycle on Medical Affairs content flat. Pharmacovigilance validation slipping on agent intake. Veeva, Salesforce, and IQVIA ship AI on their release cycle, not yours. Buy what the vendor layer can carry. Own the validation discipline that makes it ship under GxP and Sunshine Act scrutiny.

02 · PRODUCT DEFENSIBILITY

What does our commercial-tech do that the vendor AI module cannot copy for every affiliate?

Your Pre-call summary is on a vendor upgrade cycle. Your HCP interaction corpus, consent-scope graph, and MSL-to-KOL chain compound. The reasoning behind ten years of Medical Affairs engagement sits in three senior commercial-ops people, and two of them had recruiter calls last month. The moat sits underneath the vendor layer, not above it.

03 · CAPITAL ALLOCATION

Is our commercial-AI budget one instrument or two?

Plateau capital runs the Veeva, Salesforce, and IQVIA stack cheaper. Compounding capital builds the 2028 HCP-data-and-consent-graph moat. On one licence hurdle the first wins every quarter. On one scorecard the second does not exist. The CTO who walks into the GM with one budget runs the same programme every affiliate is running. The one who walks in with two, each defended, authors the decade.

What the affiliates winning this play are actually doing

Four concrete moves any commercial CTO can start this quarter. Each addresses one sub-function. Pulled together, they become one architecture the affiliate builds on.

1**Build engine: absorb vendor-AI velocity without MLR and pharmacovigilance integrity rotting**

Instrument review depth on agent-generated Medical Affairs content as a floor, not just MLR cycle time. Mandatory named-reviewer gates on every agent-generated output above a complexity threshold: no named review, no release. Validated-system discipline extended explicitly to Pre-call Agent grounding: if the agent grounds on affiliate HCP data, the grounding set is an audited artefact. Evaluation harness on Pre-call Agent and Content Agent output against affiliate HCP reality. The dashboard is no longer adoption. It is validation depth against vendor-AI throughput.

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

Three programmes that stack over eighteen months. HCP interaction corpus captured as institutional infrastructure across Vault, PromoMats, and Salesforce Life Sciences records, with consent-scope metadata per record. A consent-scope graph drawn across channels, enforceable at query time, auditable under EFPIA Code and Sunshine Act. A model-agnostic agent-memory store portable across Veeva, Salesforce, and IQVIA upgrade cycles. Integration depth graph published for Vault, PromoMats, pharmacovigilance intake, and transfers-of-value filing.

3**Product thesis: author the vendor-vs-affiliate boundary before the CAIO deck does**

Published architecture spec naming which layers the affiliate owns and which are vendor substrate. Not a slide. A document the Medical Director, Compliance Officer, and GM sign. Joint commercial-AI value document with the GM and Global Commercial CFO before any Q3 licence expansion signs. Write the cannibalisation trigger and kill criteria in Q1: if a vendor AI module starts absorbing affiliate grounding back into its cross-affiliate layer on terms the affiliate does not control, trigger the contract-renegotiation conversation.

4**Platform bench: rebuild the junior pathway around senior and agent pairing**

Your 2024 decision to stop backfilling junior commercial-ops saves Q1. It empties the 2031 architect bench. Reverse it. Juniors in 2026 do not start on routine monitoring. The agent layer absorbed that work. They start on audit of agent-generated Medical Affairs output and paired judgment work with a named senior MDM or integration lead for their first eighteen months. By 2029 you have a mid-tier bench nobody else has.

The only firm that joins authorship to working infrastructure

Your affiliate has four places to go on commercial-AI. None can do what we do.

The big life-sciences tech-transformation practices ship a standardised AI framework scaled across a pharma book. Their cost structure requires reusability. Your affiliate is idiosyncratic. The framework produces governance theatre the Medical Director rejects.

The hyperscaler pharma 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 vendor AI modules themselves have a quote on your desk. Veeva, Salesforce, IQVIA. Their product bills on licence expansion. Telling you to spend freed capacity on an affiliate-owned data moat is structurally impossible for them to say.

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

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

WHAT YOU GET

Senior advisory that speaks Medical Affairs and validation, not vendor-deck

Working infrastructure your team uses the same week

Custom agents built on your HCP corpus and your stack

The start of an HCP-data-and-consent-graph moat the affiliate owns

No vendor lock-in at any phase of the engagement

WHY IT MATTERS

Founded by operators who have carried commercial-technology seats

Business first, vendor architecture 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 commercial-tech leadership. Working infrastructure from week one.

We start with your commercial-tech leadership 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 commercial-tech function. Build the first piece of infrastructure in the room: a capital-allocation agent that takes your real commercial-AI budget, splits it two ways, and runs the alternatives analysis against your vendor roadmap. Your leadership team leaves with a shared view and the first working artefact already in hand.

WORKSHOP 02

Your commercial-AI infrastructure

A half-day build. Your Veeva lead, your Salesforce migration owner, your MDM steward, your integration lead, each leaves with five to ten pieces of working infrastructure. Validation audit agent on Medical Affairs content. Consent-scope graph prototype. HCP interaction corpus scaffolding. Integration depth map for Vault, PromoMats, pharmacovigilance intake. Infrastructure the team is using the same week.

WORKSHOP 03

Ready to run

Refined infrastructure based on real use. The document that splits the commercial-AI budget into two instruments with two hurdle rates and two kill criteria. The boundary agreement with the GM on who authors the commercial-AI thesis before the CAIO deck lands. The joint value document with the Global CFO. The eighteen-month HCP-data-and-consent-graph build plan.

Three weeks. No IT project. No Global approval needed to start. The infrastructure your commercial-tech function is already running on when we finish.

For your next commercial-tech review

The questions where your commercial-tech leadership cannot agree on an answer are the ones worth an hour on the agenda. Ask these honestly.

Is your commercial-AI budget today one instrument or two? If one, what does the GM see that splits plateau from compounding, and what is the kill criterion on each?

What grounds the Pre-call Agent summary your reps use? If it is the vendor's cross-affiliate aggregated layer, what changes in the next upgrade cycle, and who owns the authorship of that change?

If your MDM lead or two senior commercial-ops architects left for Accenture Life Sciences or Deloitte Digital tomorrow, how many months to reconstruct the HCP identity and consent-scope reasoning behind your core commercial tech?

Pre-call Agent adoption is at seventy percent. MLR cycle on Medical Affairs content has not moved. Where did the freed hours go, and is your validation bar still the one you want the Sunshine Act auditor seeing next year?

Thirty percent of European AI-strategy ownership sits with the CIO, twenty-three percent CDAO, twenty-one percent CAIO, one percent COO. Zero percent commercial CTO. Is your Q1 boundary agreement with the GM written before Global's CAIO deck, or waiting until after?

If any of these produces a disagreement in your commercial-tech leadership, that is the conversation worth having before Global 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