



AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in the Reps and MSLs You're Building

Your A-players are watching. Your MSL apprenticeship has quietly compressed. Your 2030 senior bench is being shaped by deployment decisions you were not invited to. A briefing for Chief Human Resources Officers at mid-size and specialty pharma commercial affiliates in Europe.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European specialty pharma
commercial people leaders

DATE

April 2026

Monday, 8:47 a.m.

Every CHRO at a mid-size European pharma commercial affiliate has had the same Monday. Three things in the inbox. Your Head of Talent forwarded Friday's resignation from your best hospital-specialty rep, the one three KOLs at Karolinska asked for by name. She went to a competitor pharma in the same therapeutic area, with a base twenty-five percent higher than yours. Your GM attached the Mercer thriving report to Thursday's country leadership invite, with a single-line note: "need your read before this lands with the regional HR call." Your DACH counterpart forwarded a note from her works council opening co-determination consultations on three Veeva agent deployments shipped last quarter without HR in the room.

You have held this seat for four years. The pattern used to be simple. Hire the right reps. Train them on the brand plan. Pair MSLs with senior medical affairs leads. Keep the works council in the loop. The business was understandable.

Something changed. Your call-prep time, your CRM admin time, and your first-pass MSL briefing time all dropped function-wide because Veeva, Salesforce, and IQVIA shipped agents that compress them. Your A-players are reading every deployment as a signal about whether the affiliate is being redesigned for them or around them. Your hospital-specialty rep who resigned on Friday was not the first warning. Your thriving signal dropped at the last engagement survey and nobody reconstructed the account-level story. Your works council has more evidence of shadow AI use in the affiliate than sanctioned use.

You do not have a rep problem. You have a judgment-architecture problem. Your affiliate's workforce strategy is being written every week in deployment decisions, headcount lines, and global-HQ AI playbooks you were not invited to author. The question is not whether to have one. It is whose handwriting the one you already have is in. Your A-players read it last month. Your works council is decoding it right now. Your 2030 senior MSL bench is being shaped by decisions that look like CRM rollouts and are actually architecture.

This briefing covers what is happening in your sector, what it means for the people function you run, and the three questions your country leadership team should answer before global pushes the standardised AI playbook.

The rep model just lost the equation that ran it for three decades

For three decades, the economic architecture of a pharma commercial affiliate ran on a simple exchange. Reps did call-prep, dropped samples, ran territory. MSLs read the literature, briefed KOLs, supported clinical conversations. Market access compiled HTA dossiers and tracked formulary committees. Each function had its own apprenticeship: junior reps learned territory by carrying it; junior MSLs learned the science by writing the briefing decks; junior market-access analysts learned the pharmacoeconomic models by building them. Over five to eight years, the best juniors became seniors, and the affiliate's commercial engine compounded.

The CRM-vendor AI arms race ended that exchange in eighteen months. Veeva launched Pre-call Agent, Voice Agent, and Free Text Agent through 2025-2026. Salesforce Life Sciences Cloud went general with seventy customers including Novartis, Takeda, and Pfizer. IQVIA OCE+ embedded Next Best Action. Every assumption about commercial infrastructure is temporarily unfrozen. The work that used to apprentice juniors into seniors is now what the agents do first. KOLs walk into MSL meetings having read AI-summarised trial data the MSL did not summarise. Patent clocks compress every uptake decision: four months of slower time-to-peak on a typical specialty launch is millions in gross profit that never comes back.

Your apprenticeship for senior MSLs and hospital-specialty reps ran on call-prep, briefing-deck production, and reporting work that is no longer yours to give. The question is whether the pathway gets rebuilt before the 2030 senior opening has no internal candidate.

Global HQ is preparing to push standardised AI playbooks down to affiliates within twelve to eighteen months. Once that lands, the country GM loses autonomy to shape how AI enters their organisation. Affiliate CHROs who author the redesign locally first — in one therapeutic area, with their own KOLs and their own works council — set the template global has to negotiate with. The ones who wait inherit a playbook designed in a country they have never worked in.

Your people function is feeling it every day

Your talent acquisition. Your best hospital-specialty rep resigned on Friday. Your best lateral MSL hire of the last two quarters came from an AI-native medical-information vendor who read your affiliate's hesitation as opportunity. Your medical-rep offer-accept rate is holding, but the best candidates are asking in second-round interviews what your affiliate's AI strategy is, and "global is working on it" is not an answer that wins them. A-players do not quit on the spot. They start looking quietly. By the time retention moves in the dashboard, they have already decided.

Your L&D. Your affiliate's L&D budget is allocated to compliance modules, brand-plan rollouts, and an annual sales meeting. None of it is designed for continuous in-the-work coaching, agent-paired call-prep, or deliberate cohort rotation across therapeutic areas. Two populations are forming inside your affiliate whether you designed it or not. The top decile is using AI daily and building judgment that compounds weekly. The rest is running the call-prep ritual she ran in 2022 because nobody told her how to use the new tools.

Your compensation architecture. Your rate card and bonus structure reward calls per quarter, share of voice, and territory volume. Agents are compressing the call work. Your senior MSLs and hospital-specialty reps see the math earlier than your committee does. With competitor pharmas in the same therapeutic area paying more for the same scientific credibility, the top fifteen percent of your commercial workforce are reading every quarter's compensation run for a signal about whether the affiliate is being redesigned around them or without them.

Your employee relations and works council. Your DACH works council has opened co-determination consultations on three Veeva and Salesforce deployments that shipped without HR in the room. The regulatory load on people functions has stacked: AI literacy documentation, high-risk recruiting and performance-management systems, works-council framework agreements, data-minimisation when the work product touches HCPs and patients. None of it shows up on the country GM's uptake dashboard. All of it lands on you.

Your workforce planning and succession. Your 2030 senior MSL bench has six names on it. Three came through the literature monitoring, briefing-deck production, and KOL-profile maintenance work that agents have already absorbed. You have not replaced the pathway. The external senior MSL hire in 2030 will cost a thirty to forty percent premium, take eighteen months to integrate, and bring the previous affiliate's KOL relationships rather than yours. On a long enough horizon, the CHRO who cut the junior layer in 2024 is the CHRO the regional HR call blames for the 2031 opening with no internal candidate.

The apprenticeship clock was always the architecture

Your sector is repricing from below because the call-prep and MSL briefing layer is compressing. Your people function is stretched because every Veeva or Salesforce deployment, every headcount line, every framework-agreement consultation is a fragment of a workforce strategy that somebody is writing and you are not always in the room when they do. Same force, two altitudes. But there is a third thing happening that most mid-size pharma affiliate CHROs have not named yet.

Your affiliate's workforce strategy is being written every week. The Veeva Pre-call Agent activated in field force ops. The headcount line cut in marketing PMs. The global brand-plan template that now embeds AI-augmented call planning. The framework-agreement consultation opened in DACH on three deployments. Your A-players read the handwriting last month. Your works council is decoding it right now. Your 2030 senior MSL bench is being shaped by decisions that look like CRM rollouts and are actually architecture.

The teaching hospital trains a department head in twelve years because the medical student walks the round with the consultant from week one, watching every diagnostic call, every ethical judgment, every conversation with the family. The student does not learn the science from a textbook and then meet a patient in year three. Co-location with real cases is the architecture. That architecture is available to your affiliate for the first time, because agents make it affordable to put junior reps and MSLs next to senior judgment calls again, instead of next to call-prep tasks.

You can tell which camp an affiliate is in from a single question. Of your successors on the 2030 senior MSL bench, which ones came through work that still exists? Most affiliates cannot answer. The ones that can have already rebuilt the pathway.

You have twelve to eighteen months to author the answer locally, or the global standardised playbook writes it for you in a country you have never worked in. Same therapeutic areas. Same workforce. Different outcome depending on whose handwriting the answer is in.

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

Every CHRO has had the AI conversation. Most affiliates are making the same four decisions. A smaller group is making different ones.

THE DEFAULT PATH

Restrict AI tools for compliance and safety

Sixty-nine percent of reps and MSLs use AI anyway. A-players read restriction as mistrust and start looking quietly.

Cut the junior rep and MSL layer for the headcount target

Cheaper function in Q2. Empty senior bench by 2031. External senior hire at a competitor-pharma premium the affiliate cannot absorb.

Wait for global HQ to push the standardised AI playbook

Country autonomy closes in twelve months. The redesign arrives shaped by a country you have never worked in.

Run the annual engagement survey as your main HR deliverable

Your cadence is slower than the GM's quarterly uptake review. Cost-out writes the workforce answer. You inherit the damage.

THE WINNERS' PATH

Publish three rules tied to compliance sensitivity

Promotional content stays human-authored. Non-promotional intelligence runs on agents with human review. Audit trail builds itself.

Ship one therapeutic area redesigned as a triad unit

MSL plus hospital-specialty rep plus agent ensemble. Junior sits inside every judgment call from week one. MSL apprenticeship compresses eight years to four.

Five Nordic KOLs co-author the engagement model first

One-day session. The KOLs become the authors of what reps and MSLs do with AI in the room. Global negotiates with your template.

Publish the for-them memo in first 100 days

Two pages, co-signed by the GM, referenced every quarter. Every Veeva or Salesforce deployment gets decoded against it.

Same workforce. Same twelve months. Opposite outcomes. The difference is whose hand is on the pen.

THREE QUESTIONS

Talent Gravity. Capability Compounding. Succession Readiness.

Three questions every CHRO is tracking. None of them used to be the same question. They are now.

01 · TALENT GRAVITY

Are your A-players staying because of you, or despite you?

Your best hospital-specialty rep resigned on Friday to a competitor pharma at twenty-five percent more. Your best lateral MSL last quarter came from an AI-native medical-information vendor. Your medical-rep recruiting is holding, but the best candidates are asking what your AI strategy is in second-round interviews. A-players watch which way you lean the first time something goes wrong. The ones who start looking quietly have already decided.

02 · CAPABILITY COMPOUNDING

Is the median rep sharper in December than they were in January?

Two populations are forming in every affiliate. The top decile is using AI daily and building judgment that compounds weekly. The rest is running the 2022 call-prep ritual. Your L&D budget funds compliance modules and an annual sales meeting. Neither moves the median. The affiliates pulling ahead are making the 90th-percentile rep workflow the affiliate's default workflow.

03 · SUCCESSION READINESS

Of your 2030 senior MSL bench, how many came through work that still exists?

Your successor list has six names. At least three came through the literature monitoring, briefing-deck production, and KOL-profile maintenance work that agents have already absorbed. You have not replaced the pathway. The external senior MSL hire in 2030 costs a thirty to forty percent premium and brings the previous affiliate's KOL relationships. The apprenticeship architecture gets rebuilt now or gets inherited as a gap.

What the affiliates winning this play are actually doing

Four concrete moves any CHRO can start this quarter. Each addresses one sub-function. Pulled together, they become one answer the affiliate compounds on.

1

Talent: publish the for-them memo in your first 100 days

Two pages, signed by the GM. What this affiliate believes its reps and MSLs are for. Which investments are being made in people becoming more capable. Which roles are being redesigned around judgment and agent partnership. What happens to someone whose role changes materially. Referenced every quarter. Every Veeva or Salesforce deployment after it gets decoded against this memo. Without the memo, every deployment is decoded against rumour.

2

L&D: ship one therapeutic area redesigned as a triad unit

MSL, hospital-specialty rep, and agent ensemble as one inseparable customer-facing unit on the same hospital account. The junior rep does not run separate territory. She operates the agent ensemble in front of the senior MSL on live KOL prep from week one. The MSL watches the judgment calls in real time and intervenes at the decision-level moments. Your exposure-to-judgment hours for that junior rep go from a hundred and fifty a year on the old model to twelve hundred on this one. The MSL apprenticeship clock compresses from eight years to four. Start with one therapeutic area.

3

Succession: five KOLs co-author the engagement model first

Do not roll out the AI-augmented engagement model and ask KOLs to react. Invert. Invite five Nordic KOLs from your highest-peak-share therapeutic areas into a one-day session as authors of what your affiliate's reps and MSLs do with AI in the room. They name the questions they want answered before the next visit. They name the data they want pre-read. The reps and MSLs become executors of an engagement model the KOLs helped design. Global negotiates with your template, not the reverse.

4

Trust: three rules by compliance sensitivity, one escalation, one red line

Three rules about what reps and MSLs can do with AI without asking permission. One rule about what requires medical-affairs review before it leaves the affiliate. One rule about what never goes to an agent. Promotional content stays human-authored. That is the regulatory boundary. Non-promotional intelligence runs on agents with review. Published. Defended to the GM and the works council when an experiment fails. Your A-players watch which way you lean.

The only firm that joins workforce authorship to working infrastructure

The pharma affiliate CHRO has four places to go. None can do what we do.

The pharma-specialised commercial-effectiveness consultancies (ZS Associates, IQVIA Consulting, Trinity Life Sciences, ClearView, Cello Health) sell call-plan optimisation and rep-effectiveness benchmarks. Their benchmark becomes a liability when the rep model itself is being repriced.

The CRM-vendor implementation partners (Deloitte, Accenture, Capgemini — all now Salesforce or Veeva partners) sell the platform implementation that frames the rep redesign. The framing is the problem.

The HR-tech vendors (Workday, SuccessFactors, Eightfold, Gloat) sell platform software. The data encoding which reps hold which hospital relationships IS your apprenticeship architecture in machine-readable form. Building on the vendor stack licenses it back.

The L&D platforms and pharma training academies (LinkedIn Learning, Coursera, Veeva Compass, Sana Labs) sell content libraries. The senior-MSL pathway is a pattern-exposure problem, not a content-library problem. A platform that sells modules cannot rebuild what agents have collapsed.

The market offers four things you have already said no to. GRAIL is the fifth option.

WHAT YOU GET

Senior advisory that speaks GM, not HR-deck

Working infrastructure your team uses the same week

One therapeutic area redesigned as a triad unit

A for-them memo the GM co-signs in your first 100 days

No vendor lock-in at any phase of the engagement

WHY IT MATTERS

Founded by operators who have carried commercial 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 country leadership. Working infrastructure from week one.

We start with your country 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 people function. Draft the first page of your for-them memo in the room, with your GM present. Your country leadership leaves with a shared view of what the affiliate believes its reps and MSLs are for and the first draft of the document the GM will co-sign.

WORKSHOP 02

Your people infrastructure

A half-day build. Your Head of Talent, Head of L&D, Head of Compensation, and Head of Employee Relations each leave with working infrastructure tuned to their function. Task-level role-mapping agent for the rep and MSL roles. KOL council facilitation architecture. Framework-agreement drafting for the works council. Three-rules trust document. Infrastructure your team is using the same week.

WORKSHOP 03

Ready to run

Refined infrastructure based on real use. The first therapeutic area redesigned as a triad unit with the incumbent reps and MSLs in the room. The published for-them memo, co-signed. The KOL council kickoff agenda. The GM conversation with numbers the country leadership can check: external-hire premium math, time-to-peak compression forecast, successor-readiness count.

Three weeks. No IT project. No global approval needed to start. The infrastructure your people function is already running on when we finish.

For your next country leadership team

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

If you asked three of your best reps and MSLs privately tomorrow whether the affiliate's AI strategy is being done for them or to them, what would they say? Do you know the names, and can you describe the evidence for what they would say?

Walk me through one role in your affiliate that has been redesigned at the task level in the last twelve months. Who did the redesign, what did it change, and what does the incumbent do with their week that they did not do before?

Of the six names on your 2030 senior MSL bench, how many came through the literature monitoring, briefing-deck production, and KOL-profile work that agents have already absorbed? What is replacing the pathway for the ones whose pathway is gone?

Your workforce strategy is being written this quarter in deployment decisions, headcount lines, and global-HQ AI playbooks. How many of those fragments are you the author of, and how many are you decoding after they ship?

Of your top KOLs in your highest-peak-share therapeutic areas, have any of them been asked to co-author the AI-augmented engagement model yet, or are they being presented with a model designed without their input?

If any of these produces a disagreement in your people leadership, that is the conversation worth having before global pushes the standardised playbook.

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

▶ YOU ARE READING

Capacity. Capability. Promise.

The biggest talent transition of a career. Compete on capacity, not headcount.

CPTO

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