



AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in the Bench You're Losing

Your senior class leads are five years from retirement. Your mid-tier pathway has quietly compressed. Your 2030 underwriting and actuarial bench is being shaped by decisions you were not invited to. A briefing for Chief Human Resources Officers at mid-market European specialty insurers.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market specialty
insurance people leaders

DATE

April 2026

Tuesday, 8:47 a.m.

Every CHRO at a mid-sized European specialty insurer has had the same Tuesday. Three things in the inbox before the reserve review. Your Head of Talent forwarded Friday's notice from your best mid-tier casualty underwriter, the one on the class-lead track, the one three brokers asked for by name. She went to an AI-native MGA. Your CFO attached the Mercer thriving report to Wednesday's audit-committee calendar invite, with a single-line note: "need your read before this lands with the board." Your works council opened framework-agreement consultations on three platform deployments — sanctions clearance, claims-triage, an underwriting copilot — that shipped without HR in the room.

You have held this seat for four years. The pattern used to be simple. Hire the right underwriters and actuaries. Train them well. Keep the works council in the loop. Protect the senior bench. The carrier was understandable.

Something changed. Your mid-tier pathway compressed quietly over 2024 and 2025 because the routine triage, aggregation, and standard reserving work that trained mid-tier into class leads is running on agents. Your senior class leads are reading every platform deployment as a signal about whether the carrier is being redesigned for them or around them. Your casualty underwriter who resigned on Friday was not the first warning. Your thriving signal dropped at the last engagement survey and nobody reconstructed the line-of-business story. Your works council has more evidence of shadow AI use on the desk than sanctioned use.

You do not have a workforce problem in 2030. You have an authorship problem in 2026. Your carrier's workforce strategy is being written every week in decisions you were not invited to. The question is not whether to have one. It is whose handwriting the one you already have is in. Your senior class leads read it last quarter when the platform shipped. Your works council is decoding it right now. Your 2030 underwriting bench is being shaped by IT line items that 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 leadership team should answer before the next GMT.

Specialty insurance just lost the apprenticeship that ran it for a generation

For a generation, the economic architecture of a mid-market specialty carrier ran on a simple exchange. Mid-tier underwriters did the routine triage, sanctions clearance, aggregation runs, and standard rating work under the eye of a class lead. Class leads exposed them by osmosis to the override calls — the moments where the model was wrong, the broker was packaging risk creatively, the loss triangle did not match the class. Over fifteen to twenty years the best mid-tier underwriters became class leads, and the carrier's combined ratio compounded on judgment that had never been written down. Actuarial worked the same way: trainee, fellowship, Head of Reserving. Routine work was the substrate. Override judgment was the asset.

The AI-native entrants in your sector are not experimenting. They are licensed and binding risk, with structurally lower cost bases and faster turnaround at the broker desk. The substrate work — triage, aggregation, sanctions, rate-card application — runs in seconds where your desk runs in days. The replacement structure has no mid-tier pathway: senior judgment at the top, an agent layer underneath, no bench in between. Every mid-market specialty carrier is making the same quiet decision about the mid-tier layer this year. Almost none are rebuilding the pathway.

Your apprenticeship ran for a generation on routine work that is no longer yours to give. The question is whether the pathway gets rebuilt before the 2030 class-lead opening has no internal candidate.

The retirement curve is the second weight. The senior class leads and casualty actuaries who carry your override judgment built it on twenty to thirty years of loss experience, broker-counsel relationships, and class-specific pattern recognition that has never made it into a system. The retirement window is concentrated. The combined-ratio impact of knowledge loss is invisible for twelve to eighteen months and then reads as adverse development on a book the reinsurer used to take with light scrutiny. The CEO who cannot articulate a compelling carrier-level workforce thesis risks losing the people whose departure costs more than any expense-ratio programme ever gained.

Your people function is feeling it every day

Your talent acquisition. Your best mid-tier casualty underwriter resigned on Friday. Your best lateral hire of the last two quarters came from an AI-native MGA who read your carrier's hesitation as opportunity. Your graduate and actuarial-trainee offer-accept rate is holding, but the strongest applicants are asking in second-round interviews what your AI strategy is, and "we are thinking about it" does not win 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 carrier's L&D budget is allocated to compliance modules, mandatory CII and actuarial-society programmes, and an annual conference. None of it is designed for continuous in-the-work coaching, agent-paired underwriting practice, or deliberate cohort rotation across classes. Two populations are forming inside your desk whether you designed it or not. The top decile is using AI daily and building judgment that compounds weekly. The rest is practicing not-knowing. Both groups are becoming who they practice being.

Your compensation architecture. Your rate card and class-lead compensation reward billable hours and book-size growth. Agents are compressing the substrate hours. Your senior class leads see the math earlier than your committee does. With AI-native MGAs offering retention-style packages and competitor carriers reading the demographic curve, the top fifteen percent of your senior bench reads every quarter's compensation run for a signal about whether the carrier is redesigning around them or without them. The comp committee that met last month did not see this.

Your employee relations and works council. Your works council opened framework-agreement consultations on three platform 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, framework-agreement openings on every agent-assisted deployment, the demands-and-needs paper trail under IDD, data-minimisation in claims-triage tools. None of it shows up on the cost-reduction deck on the CFO's desk. All of it lands on you.

Your workforce planning and succession. Your 2030 senior-underwriter and Chief-Actuary bench has six names on it. At least three came through the routine triage and reserving work that agents have already absorbed. You have not replaced the pathway. The external senior class lead in 2030 will cost a thirty to forty percent premium, take eighteen months to integrate into your classes, and bring the previous carrier's loss experience. On a long enough horizon, the CHRO who cut the mid-tier layer in 2024 is the CHRO the board blames for the 2031 class-lead opening with no internal candidate.

The apprenticeship clock was always the architecture

Your sector is repricing from below because the substrate underwriting layer is compressing. Your people function is stretched because every deployment decision, every IT line item, every framework agreement is a fragment of a workforce strategy that somebody is writing and you are not always in the room when they do. These look like two problems. They are one. Same force, two altitudes. There is a third thing happening that most carrier CHROs have not named yet.

Your carrier's workforce strategy is being written every week. The underwriting copilot approved by the CUO. The sanctions-clearance platform deployed in operations. The headcount line cut in claims. The framework agreement opened without HR in the room. Your senior class leads read the handwriting last quarter. Your works council is decoding it right now. Your 2030 underwriting bench is being shaped by decisions that look like IT line items and are actually architecture.

Aviation came within years of losing institutional flying judgment in the 1950s when the airframe changed faster than the apprenticeship could compound. The fix was not a course catalogue. It was a training-track redesign: every junior pilot flew with a captain on the new airframe from hour one, not after a thousand hours of solo radial-engine time. The judgment that mattered transferred forward. The judgment that did not retired with grace. That redesign is available to your carrier for the first time in a generation, because agents make it affordable to put mid-tier underwriters next to class leads on every override call from week one, instead of next to a triage queue for five years.

You can tell which camp a carrier is in from a single question. Of your successors on the 2030 bench list, which ones came through work that still exists? Most mid-market carriers cannot answer. The ones that can have already rebuilt the pathway.

You have twelve to eighteen months to author the answer deliberately, or the cost-reduction deck on the CFO's desk writes it for you. Same technology. Same workforce. Different combined ratio in 2030 depending on whose handwriting the answer is in.

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

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

THE DEFAULT PATH

Restrict AI tools for safety and compliance

The desk uses ChatGPT anyway. A-players read restriction as mistrust and start looking quietly.

Cut the mid-tier layer for this year's number

Cheaper desk in Q1. Empty class-lead bench by 2031. External senior at a premium the carrier cannot absorb.

Announce carrier-wide redesign to the GMT vote

Class leads whose comp gets compressed have the power to slow it. Vote stalls. Lift-out risk to AI-native MGAs rises.

Ship the annual engagement survey as your main HR deliverable

Your cadence is slower than the CFO's quarterly. Cost-reduction writes the workforce answer. You inherit the damage.

THE WINNERS' PATH

Publish three rules that defend real trust

Routing by risk sensitivity, not tool. Audit trail writes itself. The desk watches when an experiment fails.

Ship one underwriting role redesigned as a triad unit

Class lead + mid-tier + agent ensemble as one unit. Mid-tier sits inside every override call from week one.

Six class leads author their own class redesigns first

Eight-week council cadence. They bring it to the vote as work they authored. Lift-out risk converts to ownership.

Publish the carrier's workforce answer in first 100 days

Two pages, co-signed by CEO and CUO, referenced every quarter. Every deployment gets decoded against it.

Same desk. Same twelve months. Opposite combined-ratio outcomes by 2030. The difference is whose hand is on the pen.

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 mid-tier underwriter resigned on Friday to an AI-native MGA. Your best lateral last quarter came from the same place. Your graduate and actuarial-trainee programme is holding, but the strongest applicants are asking what your AI strategy is in second-round interviews. Senior class leads 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 underwriter sharper in December than they were in January?

Two populations are forming on every desk. The top decile is using AI daily and building judgment that compounds weekly. The rest is practicing not-knowing. Your L&D budget funds compliance modules and an annual conference. Neither moves the median. The carriers pulling ahead are making the 90th-percentile underwriter workflow the desk's default workflow.

03 · SUCCESSION READINESS

Of your 2030 class-lead bench, how many came through work that still exists?

Your successor list has six names. At least three came through the routine triage and reserving production work agents have already absorbed. You have not replaced the pathway. The external class lead in 2030 costs a thirty to forty percent premium and brings the previous carrier's loss experience. The apprenticeship architecture gets rebuilt now or gets inherited as a gap.

What the carriers 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 carrier compounds on.

1

Talent: publish the carrier's workforce answer in your first 100 days

Two pages, co-signed by the CEO and CUO. What this carrier believes its underwriters and actuaries are for. Which investments are being made in people becoming more capable. Which roles are being redesigned around override judgment and agent partnership. What happens to someone whose role changes materially. Referenced every quarter. Every deployment after it gets decoded against this memo.

2

L&D: ship one underwriting role redesigned as a triad unit this quarter

Senior class lead, mid-tier underwriter, and agent ensemble as one inseparable delivery unit. The mid-tier does not run aggregation for the class lead. The mid-tier operates the agent ensemble in front of the class lead on live broker submissions from week one. Exposure-to-override hours go from one hundred fifty a year to twelve hundred. Apprenticeship clock compresses from twenty years to nine.

3

Succession: six class leads author their own class redesigns first

Do not ship a carrier-wide redesign to the GMT vote. That vote stalls because the class leads being asked to approve it are the ones whose comp gets compressed. Invert. Assemble six senior class leads — the ones whose books are largest and whose retirement window is closest — as the authoring council. Eight-week cadence. They bring it to the vote as work they authored.

4

Trust: publish three rules, one escalation, one red line

Three rules about what the desk can do with AI without asking permission. One rule about what requires review before it leaves the carrier. One rule about what never goes to an agent. Published. Defended to the board when an experiment fails. Senior class leads watch which way you lean the first time something goes wrong. Published rules create capability. A thousand small no's create the illusion of control.

The only firm that joins workforce authorship to working infrastructure

The specialty-insurance CHRO has four places to go. None can do what we do.

The Big Four human-capital practices (Deloitte, Mercer, Korn Ferry, Willis Towers Watson) run survey-benchmark models with audit-relationship conflicts. Every carrier's answer looks different. A benchmark smooths the differences that actually matter when the workforce architecture is the asset.

The insurance-specialist HR consultancies and search firms (Howden People, NMG, Eames, IGS) know the sector. Their model is search and benchmark, not redesign. The next senior class lead retiring still walks out with the override judgment.

The HR-tech vendors (Workday, SuccessFactors, Eightfold, Gloat, Phenom, Personio) sell platform software. The data encoding which class leads carry which override patterns IS your apprenticeship architecture in machine-readable form. Building on the vendor stack licenses it back to the vendor.

The L&D platforms and exam-prep vendors (LinkedIn Learning, Coursera, Sana Labs, Josh Bersin Academy, CII online) sell content libraries and exam prep. The override judgment that determines combined ratio is a pattern-exposure problem, not a content-library problem. A platform 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 GMT, not HR-deck

Working infrastructure your desk uses the same week

One underwriting role redesigned at task level with the incumbent in the room

A workforce answer the CEO and CUO co-sign in your first 100 days

No vendor lock-in at any phase of the engagement

WHY IT MATTERS

Founded by operators who have carried people-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 people leadership. Working infrastructure from week one.

We start with your people 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 carrier's workforce answer in the room, with your CEO and CUO present. Your leadership team leaves with a shared view of what the carrier believes its underwriters and actuaries are for, and the first draft of the document the CEO 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 one underwriting class. Class-lead-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 underwriting role redesigned as a triad unit with the incumbent in the room. The published workforce answer, co-signed. The class-lead-council kickoff agenda. The CFO conversation with numbers the CEO can check: external-hire premium math, cohort-retention forecast, successor-readiness count.

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

For your next GMT

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 class leads privately tomorrow whether the carrier'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 underwriting or actuarial role in your carrier 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 class-lead and Chief-Actuary bench, how many came through the routine triage and reserving 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, IT requisitions, and framework-agreement openings. How many of those fragments are you the author of, and how many are you decoding after they ship?

Of your six most senior class leads — the ones whose retirement window is closest and whose books are largest — have any of them been asked to author the redesign yet, or are they waiting to be presented with a vote on yours?

If any of these produces a disagreement in your people leadership, that is the conversation worth having before the GMT vote.

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