

AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in the Reserves You Sign

*Your next €10M of capital. The combined ratio you print.
The embedded-value thesis your board can recite back.
Three questions that look separate, but are actually one. A
briefing for CFOs at mid-market European specialty
insurers.*

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market
specialty insurance CFOs

DATE

April 2026

Tuesday morning, reserve review before the audit committee

Every CFO at a mid-sized European specialty insurer has had the same Tuesday morning. Four things on the screen at 8:30. Q3 combined ratio printed at 96.8% against a 93% plan, reserve strengthening in motor and property doing most of the damage. Your best senior marine underwriter retired ten months ago and the adverse development in your 2023 and 2024 books is now surfacing. Your CUO wants €3M next year for an AI-native pricing engine, and his one-page says "pricing and reserving infrastructure investment." Your reinsurer's underwriting panel has asked for the loss-ratio story before December 1 renewal, and the rating-agency analyst put in a call request for Friday.

You have signed fifteen reserve numbers for this board. The pattern used to be simple. Hold the underwriting discipline, hold the reserves prudent, grow the book into the plan. Close the quarter clean. Keep the external actuary happy. The business was understandable.

Something changed. Your combined ratio is drifting and your actuarial team cannot reconstruct the line-of-business story fast enough for Friday's analyst call. Your AI-pricing capital is routing through the IT budget as a productivity line. AI-native MGAs are quoting in your specialty classes at half your turnaround, and the broker panel is putting both quotes in front of the underwriter. Your Controllers are producing Solvency II QRTs faster than they ever have, and the audit committee is asking about traceability on agentic reserving output. Your reporting is clean. The economic-capital story underneath it is not.

The job has not changed in fifty years. The bar has risen on every measure of doing it, simultaneously, and the contradictions you have been navigating since you took the seat are suddenly re-openable. The CFOs who treat them as engineering problems run the offensive play. The ones who treat them as permanent trade-offs lose capital-allocation authority to the CUO and the Chief Actuary within eighteen months.

This briefing covers what is happening in your sector, what it means for the finance function you run, and the three questions your team should answer before the next reserve review.

Specialty insurance just lost the equation that priced the sector for a decade

Every pricing model your team runs priced the book on three inputs: rate, exposure, and loss-cost multiplier. AI-native MGAs in your specialty classes are pricing on twelve to twenty inputs refreshed continuously as new submission data arrives, and they are binding risk at a third of your turnaround. Broker panels at your best accounts are now running comparison quotes where your quote sits next to theirs. You are pricing on one mechanic. They are pricing on three. Across the renewal book the mechanical difference is one to three points of loss ratio a year, and it is already showing up in the adverse-development pattern on the accounts where the comparison is clearest.

Your competitors are not waiting. Corgi is a licensed AI-native carrier. Augmented UW and Testudo are binding Lloyd's-backed risk. Traditional specialty carriers are responding by acquiring MGA capacity as tuck-ins rather than rebuilding. Some are migrating pricing on their AI-augmented lines first and letting the rest of the book follow. A small set of the more patient ones are building the underwriting-intelligence asset that AI-native entrants cannot replicate in eighteen months: captured override-rationale, loss-experience patterns from senior underwriters, fraud signals, claims-handler notes, broker-submission context. They are the quiet compounders. The rest of the market is shipping model refreshes at the same quarterly cadence and wondering why the combined ratio is drifting.

Your pricing inputs are improving. Your pricing response time is not. That is not an actuarial-model problem. It is a capital-allocation decision the finance seat has not yet claimed.

The sector is repricing underneath everyone at the same time. Broker consolidation gives fewer intermediaries more power over placement, and AI-assisted placement tools are reducing the informal advantage of long-standing broker relationships. The specialty carriers that treat this as a distribution problem rather than a pricing-and-underwriting-intelligence problem are walking into a reinsurer quarterly and a rating-agency call they will not win.

Your capital, your reserves, and your narrative are feeling it every day

Your capital allocation. Your CUO's €3M AI-pricing-engine request is sitting in the IT budget, routed through the same mechanism you use to sign off on laptops and policy-admin licences. It is not a tooling request. It is the largest single capital-allocation decision you will make this decade, and it deserves the same discipline you bring to catastrophe capacity buying or reinsurance-structure restructuring. The CFOs who ship this as a capital line to the board, with alternatives analysis against cat capacity, reinsurance efficiency, and capital return, keep the authorship for the next ten years. The ones who sign the requisition hand the authorship to the CUO and the CIO by default.

Your reserves and your forecast. The team built the annual plan on a 93% combined ratio assumption. The Q3 print came in at 96.8%. Your Chief Actuary cannot reconstruct the line-of-business story fast enough for Friday's analyst call because the override-rationale signal lives in underwriter notes, the fraud signal lives in claims-team models the actuarial team has never integrated, and the pricing-pressure signal lives in broker feedback the finance team has never had access to. The plan is not wrong because the actuarial work was lazy. It is exposed because the data architecture underneath it was built for quarterly chain-ladder reruns, not for continuous reconstruction of the loss-emergence story.

Your reinsurance and capital structure. The board voted on a reinsurance programme that assumed the same treaty shape would work year after year. It no longer does. Your AI-native competitors are running continuous-scenario models of reinsurance alternatives and restructuring each renewal. You are running the same programme with price negotiation. The capital efficiency you are leaving on the table at each December renewal is two to five points of combined ratio over a three-year cycle. Reinsurance structure is a capital-allocation question about which part of the book earns the hurdle and which does not, and it has stopped being a once-a-year treasury exercise.

Your actuarial and FP&A bench. The standard reserving triangles, IBNR pattern recognition, first-draft Solvency II QRT assembly, and ORSA drafting that used to train a junior actuary into a Senior Actuary into a Chief Actuary are now running on agents. The Q1 headcount win is real. The 2031 succession bench is empty. The finance function that cuts the junior layer without rebuilding the entry pathway has a cheaper quarter this year and no internal Chief Actuary or Head of FP&A candidate in five.

Your embedded-value thesis. Your board has asked for the strategy-day preview. Your reinsurer's panel has asked for a loss-ratio story before December renewal. Your rating-agency analyst has asked how the book stays ahead of AI-native MGA competition. Three audiences, three different documents, each drafted by a different team. None of them are anchored in the same underwriting data. The CFOs pulling ahead write the thesis once and let the three audiences draw from the same architecture. The ones behind them are editing three decks and hoping the numbers reconcile before Friday.

You do not have three scorecard problems. You have one.

Your sector is repricing from below because AI-native MGAs bring three pricing mechanics and twelve inputs to the table and the rest of the market is on one and three. Your finance function is stretched because the reserve credibility, the capital-allocation authority, and the embedded-value thesis are each being managed as a separate programme, each drawing on a different data layer, each drafted by a different team. Same force, two altitudes. But there is a third thing happening that most mid-market insurance CFOs have not named yet, and it is the one that matters most.

Capital allocation without reserve credibility is guessing. Reserve credibility without a narrative is scorekeeping. A narrative without an underwriting-intelligence asset underneath it is marketing. The finance seat that holds all three holds one integrated asset, not three separate programmes. Every capital-allocation question, every reserve-credibility gap, every narrative weakness is a symptom of the same missing layer: the underwriting-intelligence infrastructure that feeds all three from the same data, the same governance, the same discipline.

This is the Solvency-II-implementation-level decision of this decade. Every CFO in your cohort is making it right now, whether they see it or not. The ones who personally owned the 2013-2016 Solvency II implementation kept capital-allocation authority for the next decade. The ones who signed it off as "that was a regulatory project" watched the actuarial function take the board's strategic ear and the next CFO hire come from outside. The shape is the same. The asset under decision is different. The cost of getting it wrong in 2026 is compressed into eighteen months instead of five years.

You can tell which camp a carrier is in from a single question. Can your team answer any material reserve or rating-agency question by walking the same scenario against the same data in the same meeting? Most mid-market specialty carriers cannot. The ones that can have already made the shift, and they are the ones the reinsurers and rating agencies are already re-rating.

You have twelve to eighteen months before the AI-native pricing pressure shows up in the adverse-development pattern on the renewal book. The CFOs who own this decision personally run the offensive play. The ones who do not become the CFO the board replaces when the loss-ratio drift becomes a rating-agency conversation.

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

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

THE DEFAULT PATH

Route AI capital through the IT budget

Sign the CUO's requisition. Productivity line. Authorship drifts inside eighteen months.

Rerun chain-ladder quarterly on the same data

Three-week cycle. Adverse development shows up in the post-mortem. Senior retirements surface 12-18 months later.

Price on one mechanic, ship AI free

Rate $\times$ exposure $\times$ LCM. MGAs quote on three mechanics, twelve inputs. Loss ratio drifts over eighteen months.

Cut the junior layer for the Q1 win

Headcount down fifteen to twenty-five percent. Empty Chief Actuary bench in 2031. External hire at two times cost.

THE WINNERS' PATH

Route it through the reinsurance lens

One page: AI pricing against cat capacity, treaty restructure, capital return. Board sees one choice.

Rebuild reserves as a continuous view

Five scenarios in an afternoon off claims, fraud, override signals. Adverse development visible at emergence.

Migrate pricing to three mechanics

Continuous-rate, pay-per-unit-of-risk, outcome-share. Capture AI value in rate adequacy, not broker concessions.

Redesign the junior pathway around judgment

New hires pair with agents on reserve work, capital exposure, override analysis. Bench rebuilds faster than it grew.

The default path shows a cleaner Q1. The winners' path has visible friction for four quarters and compounds from month fifteen. The difference is whether the CFO is authoring the capital or administering it.

THREE QUESTIONS

Capital Allocation. Forecast Credibility. Value-Creation Narrative.

Three questions every insurance CFO is tracking. The first is the crux. The other two are how you earn the right to answer it.

01 · CAPITAL ALLOCATION

Where does the next €10M actually compound?

AI-native pricing and reserving build?
Catastrophe capacity at December renewal?
Reinsurance-structure restructuring? Adjacent-line underwriting build-out?
Capital return to the parent? Today's default: the CUO asks, the IT budget funds, the requisition gets signed, and the capital-allocation authority slides away from the seat. The question is not which tool. The question is whose decision this is.

02 · FORECAST CREDIBILITY

Why did the combined ratio print above plan, and what is the line-of-business story?

Your board voted on 93%. Your book is printing 96.8%. Your Chief Actuary cannot walk the audit committee through the line-of-business adverse-development story because the override signal sits in underwriter notes, the fraud signal sits in claims-team models, and the pricing-pressure signal sits in broker feedback the finance team has never seen. Five reserve scenarios used to take three weeks. They now run in an afternoon. The bar has risen on what credible means.

03 · VALUE-CREATION NARRATIVE

What does this book do that compounds faster than peers?

Your board wants the strategy-day preview. Your reinsurer wants the loss-ratio story. Your rating-agency analyst wants the how-we-stay-ahead answer. Three audiences, three documents, three teams, and the numbers underneath them do not reconcile. The CFOs pulling ahead write the thesis once and refresh it continuously from the same data architecture that runs the reserve view. The three audiences read excerpts of the same document.

What the carriers winning this play are actually doing

Four concrete moves any insurance CFO can start this quarter. Each addresses one lever. Pulled together, they become one underwriting-intelligence asset the carrier runs on.

1

Capital: route the AI-infrastructure decision through the reinsurance-alternatives lens

Your CUO's €3M request is a capital-allocation decision. Treat it like one. Present the alternatives on one page: AI-native pricing build, catastrophe capacity buying, reinsurance-structure restructuring, adjacent-line underwriting, capital return. ROIC-scored against the same hurdle you use for a book acquisition. Name what you are choosing not to fund to protect the thing you are funding. That is the sentence that keeps authorship at the seat.

2

Reserves: rebuild reserve adequacy as a continuous view, not a quarterly rerun

Your reserve-adequacy view should draw from live claims data, live fraud signals, live underwriter override-rationale notes, and live broker-feedback signals as one view rather than five spreadsheets stitched together. Five scenarios compress from three weeks to an afternoon. Adverse-development emergence produces the line-of-business story in the same reserve review, not the post-mortem. The reserve view stops being a deliverable on a cadence and becomes a live artefact the Chief Actuary and the CFO own together.

3

Pricing: migrate from one mechanic to three

Continuous-rate pricing on data-rich lines. Pay-per-unit-of-risk on exposure-based products. Outcome-share or performance-linked pricing where AI touches measurable claims outcome. The rate-times-exposure base holds. The three new mechanics convert AI underwriting value into rate adequacy on the classes where AI-native MGAs are already quoting. Finance authors the migration plan. Underwriting builds against it.

4

Narrative: write the embedded-value thesis once, refresh continuously

Three sentences on one page. What does this book do that compounds faster than peers? What capital gets starved to protect it? What gets fed? Refresh the thesis continuously off the same data that runs the reserve view. Reinsurer quarterly excerpts the thesis. Rating-agency prep excerpts the thesis. Strategy-day deck excerpts the thesis. Three audiences, one artefact, four registers. The CFO authors the narrative. The board recites it back.

The only firm that joins capital authorship to working infrastructure

The insurance CFO has four places to go. None can do what we do.

The big finance transformation practices will sell a twelve-to-twenty-four-month Solvency II reporting or close-acceleration programme billed by the hour. Their audit arm constrains their governance advice. Their billable-day model requires selling substrate agents commoditise. They ship tighter reporting, not authorship of the underwriting moat.

The insurance-tech platform vendors will sell a multi-year platform implementation. Your moat has to sit on carrier-owned infrastructure or it becomes a vendor asset at contract renegotiation. The CFO who builds underwriting intelligence on a Duck Creek or Guidewire stack builds it for the vendor.

The actuarial consultancies will sell reserving review, capital-model work, and pricing studies billed by the hour. They are strong at modelling. None can ship working agents your team uses Monday, or run the capital-allocation-plus-narrative conversation the seat requires.

The AI-native pricing and reserving point vendors will sell a layer where depreciation is steepest. Buying the tool does not author the capital decision. A sharper rate, the same authorship gap.

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

WHAT YOU GET

Senior advisory that speaks reserve-floor, not deck-room

Working infrastructure your team uses the same week

Custom agents built on your book, your data, your override patterns

The start of an underwriting-intelligence layer your FP&A owns

No vendor lock-in at any phase of the engagement

WHY IT MATTERS

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

We start with your finance 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 finance function. Build the first piece of infrastructure in the room: a live reserve-adequacy agent on your real book data, running against your real claims, override, and broker signals. Your leadership team leaves with a shared view of what to do next and the first piece of working infrastructure in their hands.

WORKSHOP 02

Your finance infrastructure

A half-day build. Your Chief Actuary, Head of FP&A, Treasurer, and Head of IR each leave with five to ten pieces of working infrastructure tuned to their function. Continuous reserve-adequacy agents. Reinsurance-alternatives scenario generation. SCR coverage forecasting. Rating-agency-prep narrative drafting. Infrastructure your team is using the same week, not decks about infrastructure.

WORKSHOP 03

Ready to run

Refined infrastructure based on real use. The rollout plan to the wider finance function. The capital-allocation thesis for the next board meeting. The reinsurance-structure roadmap for December renewal. And the CEO and rating-agency conversation with numbers they can check: ROIC alternatives, reserve-scenario ranges, loss-ratio sensitivity.

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

For your next reserve review

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

Your CUO wants €3M for an AI-native pricing engine. Is it in the IT budget as a productivity line, or on the capital page as one option among five? If the former, what would it take to move it before the next board cycle?

Your board voted on a 93% combined ratio plan. Your print is 96.8%. Can your Chief Actuary walk the audit committee through the line-of-business adverse-development story in the same meeting, or does it take a post-mortem? If the latter, what data architecture is missing?

Your AI-native MGA competitors are quoting on two or three pricing mechanics with twelve or more inputs. You are on one mechanic and three inputs. What is the twelve-month plan to migrate, and who on your team is authoring it?

Your reinsurer's panel, your rating-agency analyst, and your chair have each asked for a different book-performance document in the last ninety days. Are they excerpting the same thesis, or are your teams drafting three documents that do not reconcile?

Of the junior actuaries and FP&A analysts you are cutting to agents this year, how many are the people who would have made your 2031 Chief Actuary and Head of FP&A bench? What are you doing instead to build that bench?

If any of these produces a disagreement in your finance 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

▶ YOU ARE READING

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

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