

What's Changing in the Moat You Underwrite

Your technology budget is two instruments pretending to be one. Your pricing model sits above a layer that commoditises every eighteen months. Your scorecard measures system uptime and project delivery, and neither is the 2026 question. A briefing for CTOs, CPOs, and CPTOs at mid-market European specialty insurers and MGAs.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market specialty
insurance technology and product
leaders

DATE

April 2026

Monday 9:15, technology investment review

Every CTO at a mid-sized European specialty insurer has had the same Monday. Q1 combined ratio came in at ninety-three. Coding-assistant adoption across the platform team at sixty-eight percent. Cycle time on the policy-admin release train flat against last quarter. Your VP Technology opens the review with the dashboard. He is mid-sentence when your phone buzzes. It is a broker note from one of your top distribution partners: "Corgi just quoted on seventeen of the specialty risks we sent you in March. Half the bind cycle. Sharper price." Your VP Technology mentioned a recruiter call from a London AI-native MGA over coffee Thursday. Your CEO has a Chief AI Officer shortlist from the board on his desk. Your CFO forwarded an Akur8 demo this morning: "Worth discussing before the renewals window."

You have run this function for ten years. The pattern used to be simple. Hold the policy-admin platform stable. Ship the next regulatory release on time. Modernise the underwriting workbench in eighteen-month phases. Meet the board on system uptime and project delivery. The business was understandable.

Something changed. AI-native MGAs are licensed and binding risk in your specialty classes. Corgi raised over a hundred million by January. Augmented UW is winning smart-follow placements. Testudo is Lloyd's-backed. They quote at faster turnaround on lower cost bases. Coding assistants compressed engineering throughput across your platform team. Standard claims-triage and routine actuarial reporting compressed alongside. Ninety-five percent of corporate AI deployments produce no measurable P&L impact, and the gap between what the tools can do and what your firm absorbs is where your seat now lives or loses.

You are not running one technology function. You are running two, and only one is on your scorecard. One funds what already runs: policy-admin, claims operations, regulatory reporting. The other funds what has to exist by 2028: the broker-data rights you have not renegotiated, the override-rationale corpus you have not built, the broker-integration depth that outlasts three pricing-vendor swaps, the loss-experience graph sitting in three retiring senior underwriters' heads. 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 technology and product function you run, and the three questions your team should answer before the next investment review.

Specialty insurance just lost the equation that priced the category

The AI-native MGAs in your category are not experimenting. They are bound. Corgi got regulatory approval in July 2025 and crossed forty million in ARR on a hundred and eight million raised, all inside eighteen months. Augmented UW launched in July 2025 for the smart-follow market. Testudo launched January 2026 with Lloyd's backing. InsurTech investment was up nearly twenty percent in 2025 to over five billion. Eighty-six percent of insurers plan increased AI spend in 2026. Your top broker partner already sends submissions to two of these names alongside yours, and the procurement teams are running comparison quotes where your bind cycle and your pricing sit next to the AI-native's. They are not asking which is the better carrier. They are asking which model their treasury can defend.

The moat your firm ran on for a decade has moved. The pricing model is what every AI-native already has, off the shelf. Akur8, Sixfold, and Cytora ship pricing-as-a-service. The thing that does not commoditise is the layer underneath: thirty years of bound-risk loss data, the override-rationale archive, the broker-integration depth published for your top thirty distribution partners. Those compound. They take years to build and they cannot be bought. An AI-native cannot ship them because they do not have the customer base or the broker network that generates the signals.

Your pricing model shipped on a vendor API is on an eighteen-month obsolescence curve. Your loss-experience graph and broker-integration depth are the assets the AI-native MGA cannot replicate without your customer base.

The third force is the underwriting bench. Your senior underwriters carry decades of loss experience that determines combined ratio. They override the model when the data smells wrong. They know when a market is turning. Most of that judgment has never been written down. Three of them retire inside the next five years. The combined-ratio impact of knowledge loss is delayed twelve to eighteen months, then it lands. The carriers that are quietly encoding override-rationale before retirement are building infrastructure. The ones that are still running coding assistants and calling it an AI strategy are the ones whose 2028 loss ratio is already being set.

Your technology function is feeling it every day

Your build engine. Coding-assistant adoption is at sixty-eight percent. Cycle time on the policy-admin release train flat. Velocity is up double-digit on the dashboard and people are pushing large changes barely reviewed. The foundation rots while DORA resilience requirements harden. Twelve months in, your team cannot ship a category response when an AI-native MGA arrives. The review discipline that worked in 2022 is training the wrong practice on agent-generated volume now.

Your platform and data. The platform discipline that lets agents ship cleanly into a regulated, capital-disciplined system took the firms that have it four years to build. You cannot replicate four years in eighteen months. You can buy the build-side discipline from the policy-admin AI modules, the pricing-tool layer, and the claims-AI vendors. Redirect the compounding capital into the layer the vendors cannot sell: broker-data rights renegotiated across the next twenty-four months of treaty and distribution renewals, a structured override-rationale corpus, an integration-depth graph published for your top thirty broker partners. The pricing-model layer commoditises. The data and integration layer compounds.

Your product thesis. Your board asked for the firm's three-sentence thesis last quarter. Your answer was the one from three years ago. Your rate card assumes AI-priced specialty lines will price like legacy lines. They do not. Corgi runs outcomes-priced binds. Augmented UW runs smart-follow at fractional cost. You run one mechanic. That is a capital-allocation question about where the moat gets built, not a pricing one. Your CUO and CFO need a joint document with you before the AI-native MGA forces the conversation on the broker's terms.

Your R&D bench. The senior platform architect who carries the policy-admin reasoning is fifty-eight. The senior actuary who carries reserving judgment under regime change is fifty-six. Three of your senior underwriters retire inside five years. The routine work that used to train juniors into mid-tier into seniors is now running on agents. No routine work, no training ground, no senior bench in 2030. OpenAI and Anthropic both 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 bench in five years.

Your board-level authority. Heidrick's February 2026 data on European AI-strategy ownership: thirty percent CIO, twenty-three percent CDAO, twenty-one percent CAIO, one percent COO. CTO and CPO are not on the list. Not because technology leaders abandoned the mandate. You have not yet claimed it. Your CEO has a Chief AI Officer shortlist. Every month without a visible claim, the thesis drifts further from the seat that can defend it. Your Q1 boundary agreement with the CEO and the CUO lands before the shortlist, not after.

The moat that matters in 2028 is not above the pricing model

Your sector is repricing from below because AI-native MGAs are binding outcomes-priced specialty risks at a fraction of your unit economics. Your function is compressing because coding assistants and claims-AI absorbed the routine work, and the review bar that worked for the old volume cannot hold what the platform team produces with agents. Same force, two altitudes. There is a third thing happening that most mid-market specialty CTOs have not named yet.

The moat that matters in 2028 is not above the pricing-model layer. It is the loss-experience and broker-integration graph underneath it. The pricing model your team ships next quarter on a vendor API is on an eighteen-month obsolescence curve. The next pricing platform rolls in 2027 and your model commoditises. What does not commoditise: the bound-risk loss data you have rights to, the override-rationale corpus encoding senior-underwriting judgment, the integration depth published for your top broker partners and reinsurance counterparties. A Corgi or Augmented UW cannot build that layer because they do not have the customer base or the broker network that generates the signal. You do. The window closes in roughly eighteen months, the time between now and when the AI-native MGAs have enough specialty-class footprint of their own to start building the same layer.

In 2010, a mid-tier European bank ran the best policy-admin equivalent of its day: a clean mainframe core. Uptime was nine fives. Regulatory reporting on time. The board approved another five-year mainframe maintenance plan and another phased upgrade. Klarna and Revolut and N26 and Wise built modern data-and-API platforms from day one. By 2018, the bank had lost the customer-acquisition narrative across European retail banking. The 2010 scorecard never turned red. It just stopped being the scorecard the market was watching. The job had not changed. The unit of judgment had.

You can tell which camp a firm is in from a single question. Is the technology budget you defended at the last investment review a plateau budget, a compounding budget, or one budget pretending to be one when it is two? Most mid-market specialty incumbents cannot answer. The ones that can have already made the shift.

You have twelve to eighteen months. The CTOs who claim authorship of this decision in Q1 author the decade. The ones who wait until the Chief AI Officer hire arrives inherit a seat that has narrowed to platform uptime and regulatory delivery, and a scorecard that stopped mattering the quarter their biggest broker sent the first submission to a Corgi.

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

Every CTO at a specialty insurer has had the AI conversation. Most are making the same four decisions. A smaller group is making different ones.

THE DEFAULT PATH	THE WINNERS' PATH
Roll out coding assistants, track adoption Sixty-eight percent adoption. Cycle time flat. Large changes pushed barely reviewed. Foundation rots under a green dashboard.	Buy the tool layer, redesign the review bar Review-time floor per change. ADRs mandatory on agent-generated changes. Eval harness as first-class infrastructure.
Run AI budget through IT, sign the requisition CIO owns the roadmap inside eighteen months. CAIO arrives in Q3. CTO seat narrows to platform uptime and regulatory delivery.	Route technology capital through the investment hurdle Two budgets, two hurdle rates, two kill criteria. Plateau delivers margin. Compounding delivers the 2028 moat.
License a vendor pricing model, ship to the legacy book Same rate card. Same bind cycle. Corgi quotes at half the turnaround on the same broker's submission.	Stand one outcomes-priced specialty line on protected P&L Parallel P&L and treaty. Joint Pricing Architecture Document with CUO and CFO. Seat-licensed core book holds. Outcomes-priced wins net-new.
Cut the junior layer for the Q1 headcount win Cheaper function this quarter. Empty succession bench by 2031. External CTO hire at a premium the firm cannot absorb.	Pair juniors with seniors on judgment work OpenAI and Anthropic resumed junior hiring in late 2025. Eighteen-month apprenticeship on override-judgment, not five years on routine.

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

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

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

01 · BUILD VELOCITY

Is our engineering speed shipping production-grade output, or releases that fail under DORA?

Coding-assistant adoption at sixty-eight percent. Cycle time flat. Large changes pushed barely reviewed. The platform discipline that lets agents ship cleanly into a regulated, capital-disciplined system took four years to build for the firms that have it. You cannot replicate four years in eighteen months. Buy what the tool layer can carry. Own the discipline that makes it ship.

02 · PRODUCT DEFENSIBILITY

What does our firm do that a Corgi at lower unit economics cannot copy?

Your pricing model commoditises every eighteen months when the next vendor API rolls. Your bound-risk loss data and broker-integration depth compound. Your override-rationale sits in three retiring underwriters' heads, and the combined-ratio impact of losing it lands twelve to eighteen months after they walk. The moat sits underneath the pricing model, not above it. The window to build it is eighteen months.

03 · R&D CAPITAL ALLOCATION

Is our technology budget one instrument or two?

Plateau capital runs the policy-admin platform at lower unit cost. Compounding capital builds the 2028 moat. On one hurdle rate the first wins every quarter. On one scorecard the second does not exist. The CTO who walks into the investment review with one budget runs the same programme every peer is running. The one who walks in with two, each defended separately, authors the decade.

What the carriers winning this play are actually doing

Four concrete moves any specialty insurance CTO can start this quarter. Each addresses one sub-function. Pulled together, they become one architecture.

1 **Build engine: absorb agent-generated velocity without the platform foundation rotting**

Instrument review-time-per-change as a floor, not just cycle time. Large changes require two senior reviewers and documented review. Architectural decision records mandatory on every agent-generated change above a complexity threshold: no ADR, no merge. Evaluation harness coverage as a first-class metric across pricing, reserving, and claims-triage models. Rebalance the senior-to-junior ratio around the composite shape rather than preserving the 2022 team. The dashboard is no longer velocity. It is review depth against output volume.

2 **Platform and data: build the moat underneath the pricing-model layer**

Three concrete programmes that stack over eighteen months. Override-rationale capture programme with your senior underwriters before retirement, structured so the judgment is portable across pricing-vendor swaps. A loss-experience graph maintained as institutional infrastructure, with broker-data rights renegotiated across your next twenty-four months of treaty and distribution renewals. An integration-depth layer published for your top thirty broker partners and your reinsurance counterparties. The pricing model commoditises. This layer compounds.

3 **Product thesis: stand one outcomes-priced specialty line on protected P&L**

Do not migrate the whole book. Keep the seat-licensed core for the installed broker network that does not want to switch. Stand one outcomes-priced or per-bind-risk specialty line on a separate P&L and reinsurance treaty for net-new and for brokers explicitly asking. Write the cannibalisation trigger and kill criteria in Q1: if outcomes revenue on a class exceeds seat-licensed revenue, trigger the migration conversation. The joint Pricing Architecture Document with your CUO and CFO is signed before the first outcomes-priced bind closes.

4 **R&D bench: rebuild the junior pathway around senior and agent pairing**

Your decision to stop hiring juniors saves Q1. It empties the 2031 senior bench across both technology and underwriting. Reverse it. Juniors in 2026 do not start on routine code or routine submissions. The agent layer absorbed that work. They start on audit, exception-handling, and paired composite work with a named senior for their first eighteen months. The career curve compresses because the senior's judgment is now legible in the agent-and-junior review loop. By 2029 you have a mid-tier bench the competition cannot hire away because it did not exist anywhere else to hire from.

The only firm that joins R&D authorship to working infrastructure

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

The big insurance-practice firms ship a standardised AI-governance framework scaled across a sector book. Their cost structure requires reusability. Your specialty-class seat is idiosyncratic. A reusable framework produces governance theatre you cannot defend to the board.

The policy-admin vendor pro-serve arms (Guidewire, Duck Creek, INSTANDA, Sapiens) have a quote on your desk. Their business model routes every build-or-buy decision through their roadmap. They cannot tell you to spend compounding capital outside their platform.

The pricing-as-a-service vendors (Akur8, Sixfold, Cytora, claims-AI) bill on bound-risk volume. Their product is a pricing-model intervention. They cannot tell you to spend freed capacity on the override-rationale corpus instead.

An in-house or fractional Chief AI Officer splits the seat. Strategy CAIO is revenue-measured; Platform CAIO is delivery-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 capital-allocation conversation the seat requires. GRAIL is the fifth.

WHAT YOU GET

Senior advisory that speaks engineering-review, not transformation-deck

Working infrastructure your team uses the same week

Custom agents built on your platform, your loss data, your brokers

The start of a loss-experience and broker-integration moat your firm 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 technology leadership. Working infrastructure from week one.

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

WORKSHOP 02

Your technology infrastructure

A half-day build. Your VP Technology, Head of Underwriting Platforms, Head of Claims Tech, and Head of Actuarial Systems each leave with five to ten pieces of working infrastructure tuned to their function. Review-bar audit agent. Architectural decision scaffolding. Override-rationale capture protocol. Broker-integration depth prototype. 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 technology budget into two instruments with two hurdle rates and two kill criteria. The boundary agreement with the CEO and CUO on who authors the AI thesis. The joint Pricing Architecture Document with your CUO and CFO. The eighteen-month moat build plan.

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

For your next technology investment review

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

Is your technology 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-native MGA in your specialty class. What is their pricing mechanic, and at what bind-cycle and unit economics are they quoting against you? If you cannot name the top three, that is the first thing the briefing changes.

If your three most senior underwriters retired in the next eighteen months, how many months would it take to reconstruct the override-rationale that held your combined ratio together?

Coding-assistant adoption is at sixty-eight percent. Cycle time has not moved. Where did the freed hours go, and is your review bar still the one you want the policy-admin foundation trained on?

Heidrick's February 2026 data: thirty percent CIO, twenty-three percent CDAO, twenty-one percent CAIO, one percent COO, near zero CTO-CPO on the AI-strategy-ownership list. Is your Q1 boundary agreement with the CEO and CUO written, or is it waiting until the Chief AI Officer shortlist lands?

If any of these produces a disagreement in your technology 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