Openkoda

Modernizing P&C Insurance Software

Best Strategies, Solutions and a Low-Risk Roadmap

- ✓ Vendor-Neutral Modernization Guide
- ✓ Five Strategies Compared
- ✓ Backed by 2025 and 2026 Research

A mini-ebook by Openkoda

Contents

WHO THIS EBOOK IS FOR

CIOs, COOs, CTOs and heads of product, underwriting and claims at P&C carriers, mutuals, MGAs and insurtechs who run on core systems that have become hard to change.

What you will take away: a clear picture of what legacy really costs a P&C insurer and where it hurts most, five modernization strategies compared side by side, a low-risk roadmap that replaces the core one product at a time, and eight questions for choosing a modern policy administration system.

Authors

Arkadiusz Drysch

CEO & Co-Founder, Openkoda

Michał Głomba

CRO & Co-Founder, Openkoda

Arkadiusz Krysik

Marketing Manager at Openkoda

CHAPTER 1

A Good Year Is the Best Time to Fix the Roof

For US property and casualty insurers, 2025 was the best underwriting year in a decade. AM Best puts the industry's combined ratio at 92.2, down from 95.9 in 2024 and 101.9 in 2023, the year the industry lost $21.3 billion on underwriting. In 2025 it booked an underwriting gain of **$60.9 billion**.

It is tempting to read that as a sign the hard work is done. It is closer to the opposite. Much of the recovery came from years of rate increases, and rate is already turning. Marsh's global commercial insurance index fell 6% in the second quarter of 2026, the eighth quarterly decline in a row, with property rates down 12%. Swiss Re Institute expects real non-life premium growth to slow to just 0.6% in 2026 and non-life return on equity to fall from about 14% in 2025 to 11.4% in 2026 and 7.7% by 2028.

FIGURE 1

US P&C combined ratio, 2023 to 2025

Losses and expenses as a share of earned premium; below 100 means an underwriting profit

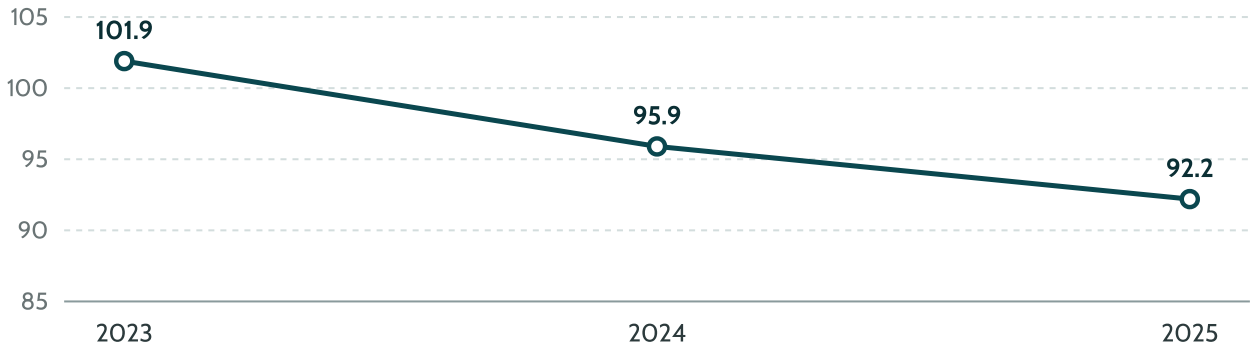

Source: AM Best, 2026 (2024 and 2025); AM Best via Insurance Journal, 2025 (2023)

Underneath the headline, loss costs keep climbing. Insured catastrophe losses topped $100 billion for the sixth year running in 2025, and Swiss Re attributes a record 92% of them to so-called secondary perils such as severe convective storms and wildfires. On the liability side, a 2025 study by the Casualty Actuarial Society and Triple-I found that legal system abuse added between **$231.6 billion and $281.2 billion** to US liability losses over ten years, up to 30.8% of booked losses in commercial auto. In motor, CCC reports that bodily injury severity rose 10.3% in a single year and 32% over four, while total losses reached a record 23.1% of claims.

FIGURE 2

Global commercial insurance rate change by line, Q2 2026

Year-on-year change in renewal rates

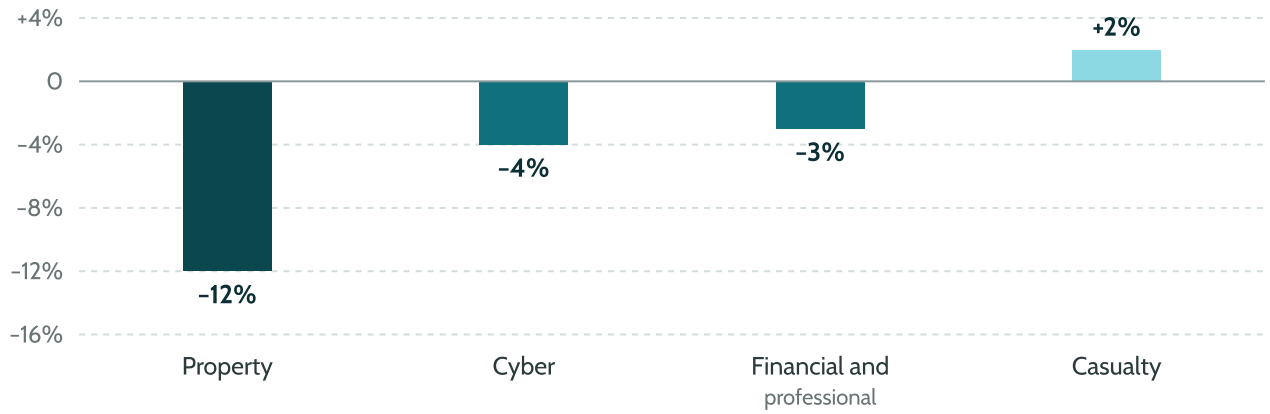

Source: Marsh Global Insurance Market Index, Q2 2026

When pricing power fades, the insurers that keep their margins are the ones that run leaner, price more precisely, settle claims faster and launch the products their customers are already asking for. Every one of those depends on software, and customers are not waiting. J.D. Power found that 57% of US auto customers shopped for a new policy in the past year, the highest level it has recorded, and that 47% of new policies are now bought through digital channels. In claims, 52% of customers who rate their insurer's digital experience as poor or just OK say they are likely to switch, against 4% of those who rate it excellent.

That is the argument of this ebook. The good years are the moment to modernize, while there is margin to invest and before the softer market makes every inefficiency visible. The chapters that follow look at what legacy really costs a P&C insurer, where it hurts most, which modernization strategies actually work, and how to make the move without the pain the industry has learned to fear.

CHAPTER 2

What Legacy Really Means in P&C

A system isn't legacy because it is old. It is legacy when changing it costs more than the change is worth. By that definition, most P&C insurers run on legacy somewhere, often in the places that matter most.

The typical stack grew one decision at a time. A policy administration system bought fifteen or twenty years ago and customized ever since. An older book still on a mainframe. Separate billing and claims systems, rating logic in spreadsheets, and a web of overnight batch files holding it all together. West Monroe's 2025 survey of 300 US insurance executives found that nearly half run core platforms six to ten years old, and some rely on systems more than fifteen years old. Adacta's European research puts the typical legacy core at 13 to 15 years.

The price of standing still

The cost shows up first in the budget. In the same West Monroe survey, **54% of insurers spend more than half of their IT budget simply maintaining existing systems**, and 52% have delayed or cancelled two or three major technology projects because the money wasn't there. Research by RSM US for INTX Insurance Software puts the hidden cost of legacy at up to $5 million a year for a single insurer: roughly 900 hours and $450,000 lost to IT support, $475,000 to $1.1 million in manual labor, and more than $1 million lost to data that arrives too late to act on. More than half of policy workflows still need a person to intervene, and 72% of insurers rely on Excel or home-grown tools to fill the gaps.

FIGURE 3

How legacy systems hold US insurers back

Share of insurance executives, 2025

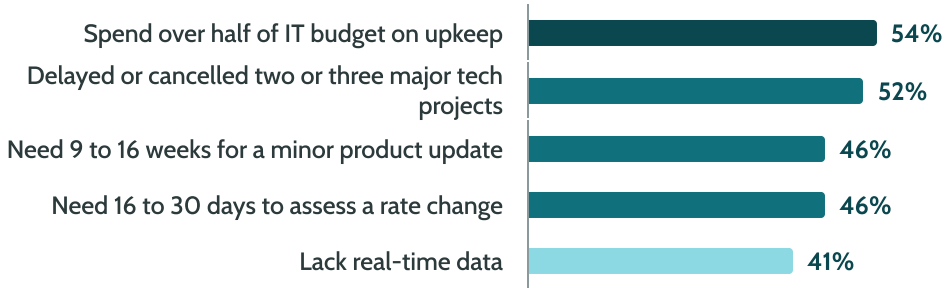

Source: West Monroe, 2025 (300 US insurance executives), via Insurance Business

Spending is rising anyway. Gartner forecasts insurance IT spending to grow 9.4% in 2026 to $256.8 billion, and Datos Insights estimates that insurers now spend 4.6% of direct written premium on IT, against about 3.7% a decade ago. A large share of that goes to the core itself: at large insurers, core systems account for more than half of the IT budget. The industry is paying more every year for the privilege of standing still.

The knowledge problem

Then there is the knowledge problem. Legacy systems are usually documented best in the heads of the people who built them, and those people are retiring. Kyndryl's 2025 mainframe survey found that 70% of organizations struggle to find the talent to modernize their mainframe estates. Every year of delay adds customizations that someone will one day have to understand, test and move. McKinsey's research with CIOs put technical debt at 20% to 40% of the value of the entire technology estate, and found that 10% to 20% of the budget meant for new products is quietly diverted to dealing with it.

WHAT THIS MEANS FOR INSURERS

The real cost of legacy is not the maintenance line. It is the product that never launched, the rate change that took a month and the claim that took six weeks. Measure those, and the business case usually writes itself.

CHAPTER 3

Where Legacy Hurts, Function by Function

Legacy rarely fails in one dramatic outage. It fails a little every day, in different ways in different parts of the business. A useful way to build the case for change is to walk the P&C value chain and ask, at each step, what the current system makes harder than it should be. The sections that follow look at each area in turn, and the summary table sets out the overall pattern.

Policy administration, products and rating

This is where rigidity costs the most, because every new product, coverage and rate change passes through it. In West Monroe's survey, 46% of insurers said a **minor product update takes nine to sixteen weeks**, and 46% need 16 to 30 days just to assess a rate change. BriteCore's annual survey of North American P&C leaders shows the gap growing: 75% say launching new products quickly is highly important, but only 42% are satisfied with how well their system delivers it, a 33-point gap that was 27 points a year earlier. When products are hard-coded, every idea becomes an IT ticket and every rate change waits in a release queue.

Underwriting

Underwriters are among an insurer's most expensive people, and legacy keeps them busy with the wrong work. Accenture's 2025 underwriting research found that underwriters in both commercial and personal lines still spend **35% of their time on non-core activities** such as re-keying data, chasing documents and switching between systems. That is an improvement on three years earlier, but a third of an underwriting team's capacity is still lost to the tools around it. Submissions arrive by email, get copied into one system for rating and another for the policy, and the reasons behind a referral live in someone's inbox.

TABLE 1

Legacy and modern, function by function

Area	What legacy looks like	What modern looks like
Products and rating	Hard-coded products; rate changes wait for IT releases	Products, rates and rules configured and versioned by the business, tested before release
Underwriting	Re-keyed submissions; a third of time on non-core work	A prefilled workbench with referral reasons and the full rating breakdown
Billing	A separate system reconciled overnight	Instalments and payments on the policy record, in real time
FNOL and claims	Phone and email intake; manual triage and updates	Digital FNOL, automated triage and proactive status updates
Distribution	Each portal or partner is a new project	Embeddable forms and APIs sharing one set of rules
Integrations and data	Batch files, point-to-point links, silos	API-first, real-time data, one source of truth
Compliance	Audit evidence assembled by hand	An automatic audit trail and fine-grained roles
AI	Data locked in old schemas; pilots that never scale	Clean, structured data and AI with human review

Billing and payments

Billing often lives in its own system, reconciled with the policy record overnight. The result is familiar to every service team: a customer calls about an instalment, and the answer depends on which screen the agent opens. Payment preferences have moved to cards, digital wallets and autopay far faster than many billing engines have.

FNOL and claims

Claims is where customers find out what their policy is really worth, and speed is the whole story. J.D. Power's 2026 property claims study found that it still takes an average of **40.7 days from first notice of loss to final payment**, and only 38% of customers reported their claim digitally. In the digital

claims study, just 22% of customers said they received adequate status updates, even though updates are one of the strongest drivers of satisfaction. Legacy claims systems make this worse: FNOL arrives by phone and email, triage is manual, and status updates depend on an adjuster finding time to send them.

FIGURE 4

What drives satisfaction in property claims

Customer satisfaction on a 1,000-point scale

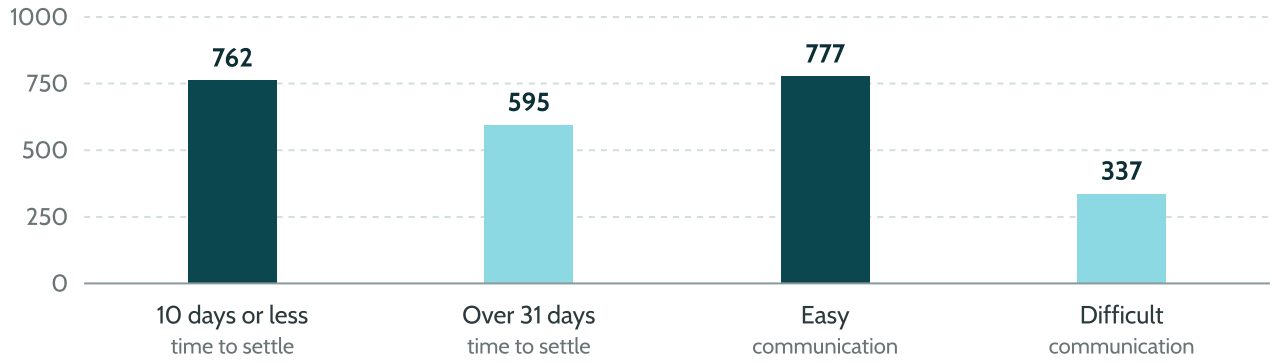

Source: J.D. Power 2025 US Property Claims Satisfaction Study, via Insurance Journal

Distribution, portals and partners

Brokers and customers now expect to quote and buy wherever they are, and new channels keep appearing. BCG expects embedded P&C premium to grow from about $13 billion today to more than $70 billion by 2030, and names **legacy systems, siloed data and rigid product structures** as what holds incumbents back. On a legacy stack, every new portal or partner is its own integration project, with its own copy of the rules.

Integrations and data

Most insurers don't struggle with any single system as much as with the connections between them. In Adacta's 2025 research, 45.5% of insurers named integration with new technology as a major limitation of their legacy core, just behind inflexibility at 46.4%. A 2025 survey of insurance leaders for SAS found that poor data quality is the biggest obstacle to good decisions (41%), and 38% lack a real-time view of their risks, revenue and costs. When data moves in nightly batches between systems that each hold part of the truth, reporting becomes an exercise in reconciliation.

Compliance, security and reporting

Regulators increasingly expect insurers to prove what their systems did, not just describe it. In the EU, the Digital Operational Resilience Act has applied since January 2025 and requires a register of every ICT third-party arrangement. In the US, 28 jurisdictions had adopted the NAIC Insurance Data Security Model Law by August 2025, according to the NAIC. Unsupported software and audit evidence assembled by hand make both harder, and a breach is expensive: IBM puts the global average cost at $4.44 million.

AI readiness

Finally, legacy decides how much value an insurer can get from AI. BCG found that only 7% of insurers have scaled AI across the business, and Capgemini's 2026 World P&C Insurance Report puts the share with advanced AI maturity at 10%. The barriers are the same ones legacy creates: in AM Best's survey of insurers, the top obstacles were data readiness, security and privacy, and integration with legacy systems.

FIGURE 5

Top barriers to AI adoption named by insurers

Share of insurers naming each barrier

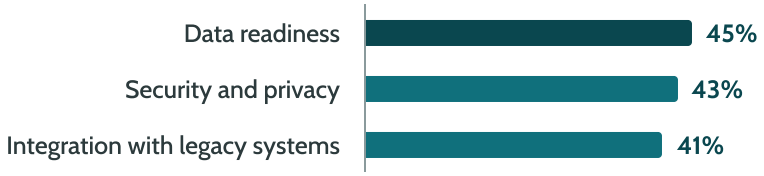

Source: AM Best survey of insurers on AI adoption, via AM Best News

CHAPTER 4

Why Insurers Wait, and Why Waiting Costs More

If legacy is this costly, why do so many insurers keep it? Because the fear of replacing it is rational. The industry has watched too many core programs go wrong to treat modernization lightly.

The numbers behind that fear are real. Roland Berger's 2026 study of core insurance transformations found that **around half of core system replacement programs miss their time or budget targets**, on programs that typically run two to five years and cost from €20 million to more than €100 million. RSM's 2026 research found that 45% of insurers had core implementation cycles longer than

18 months, and average satisfaction with those implementations came in below 72%. Across industries, the classic McKinsey and Oxford study of large IT projects found they ran 45% over budget on average and delivered 56% less value than planned.

Why core programs really fail

Look closer, though, and the failures have less to do with technology than with the way programs are run. Roland Berger names underestimated data migration, weak ownership and unclear decision-making as the main causes, and concludes that execution discipline matters more than platform selection. The same patterns come up again and again:

The big bang. Every product, every book and every process moves on one cutover date, so every risk is concentrated in a single weekend.

Rebuilding the past. Requirements are written from the old system, so the new one ends up recreating decades of workarounds instead of retiring them.

Customizing instead of configuring. Each change that goes into code instead of configuration has to be retested at the next upgrade, until the new system is legacy again.

Data left until last. Migration is treated as a technical task at the end rather than a business decision at the start, and quality problems surface when there is no time to fix them.

FIGURE 6

Where the hard part of scaling AI lies

Share of challenges insurers face, by source

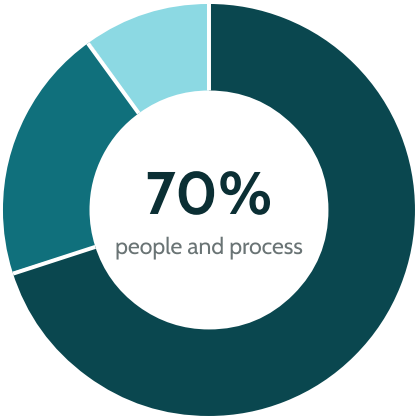

- **70%** people and processes
- **20%** technology and data
- **10%** the algorithms themselves

Source: BCG, Insurance Leads in AI Adoption, 2025

BCG sees the same thing in AI programs: around 70% of the challenges insurers face in scaling AI come from people and processes, 20% from technology and data, and only 10% from the algorithms themselves. Modernization is no different. **The technology is rarely the scary part.** The way the change is structured is what decides whether it hurts.

The cost of waiting

Waiting has costs of its own. Every year on a legacy core adds maintenance spend, customizations and dependence on a shrinking pool of people who understand the system. BCG estimates that modernizing the core costs 2% to 4% of non-life gross written premium in capital expenditure. That is a significant number, but it is paid once. The costs described in chapter 2 are paid every year, and they grow.

The market seems to have reached the same conclusion. In BearingPoint's European core systems research, the share of insurers who need a new core but are not ready for it fell from 24% to 14% in a year, while the share actively planning a new solution rose from 21% to 35%. In the US, Datos Insights reports that about two-thirds of large carriers are planning or continuing a policy administration replacement in 2026.

FIGURE 7

Where European insurers stand on core systems

Share of insurers in each stage

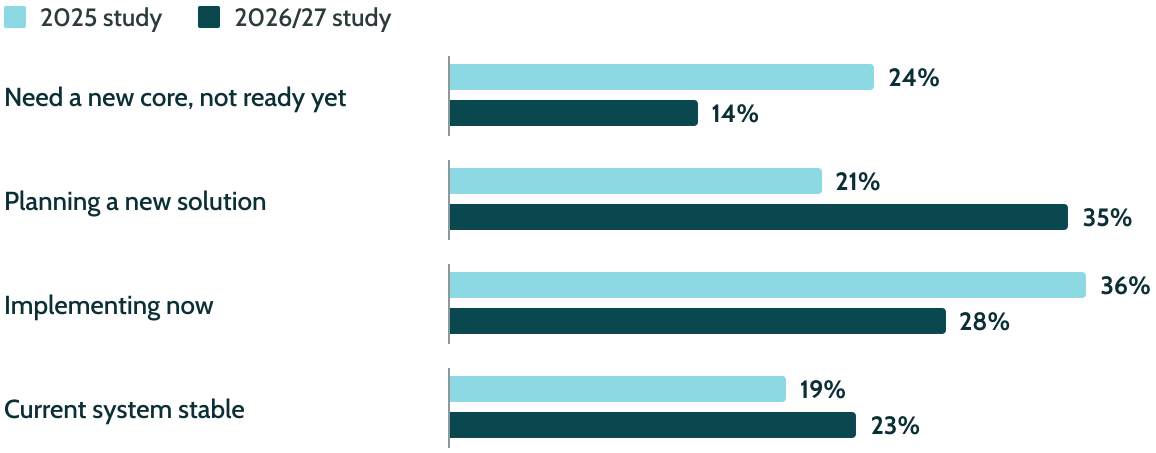

Source: BearingPoint Kaleidoscope P&C Core Systems Study, 2025 and 2026/27 editions

WHAT THIS MEANS FOR INSURERS

The question is no longer whether to modernize, but how small, safe and fast the first step can be. The rest of this ebook is about making that step.

CHAPTER 5

Five Modernization Strategies, Compared

There is no single way to modernize a P&C core. There are five broad strategies, and most successful programs combine two or three of them. What separates them is how much they change, how quickly they deliver value, and whether they fix the underlying problem or only move it.

1. Rehost or replatform. Move the existing system to the cloud or to newer infrastructure, largely as it is. It lowers hosting costs and removes some hardware risk, but the rigidity comes along for the ride. A product that took twelve weeks to change on premises still takes twelve weeks in the cloud.

2. Wrap and extend. Leave the core in place and put APIs and modern digital front ends around it: a broker portal, a customer app, a claims intake form. It is fast and low-risk, and it improves the experience people see. But every new product still has to be built in the old core, so the bottleneck stays where it was.

3. Hollow out the core. Move capabilities out of the legacy system one at a time, such as rating, underwriting rules, documents or reporting, into modern components connected by APIs. Over time, the old core shrinks to a record keeper. It spreads the risk, but it can take years, and every extracted piece needs its own integration.

4. Progressive replacement. Launch new products and programs on a modern policy administration system from day one, and move existing books across at renewal, line by line, while the legacy system keeps running what hasn't moved yet. Each step delivers value on its own and can be tested in production before the next one starts.

5. Full replacement. Replace everything in one program, with one cutover. When it works, it fixes the root problem in one go. When it doesn't, it is the source of the statistics in chapter 4.

TABLE 2

Five strategies at a glance

Strategy	What changes	Risk	First value	Fixes the root problem?
Rehost or replatform	Where the system runs	Low	Months	No, the rigidity moves with it
Wrap and extend	Front ends and APIs around the core	Low	Weeks to months	Partly; products still live in the old core
Hollow out the core	Capabilities extracted one by one	Medium	Months	Partly, over several years
Progressive replacement	New products on a new PAS; books move at renewal	Contained per product	Weeks to months	Yes, one line at a time
Full replacement	Everything, at one cutover	High	Years	Yes, if it lands

What the experts recommend

The direction of expert advice is clear. McKinsey recommends a hybrid path for most insurers: write new policies on the new platform while legacy policies stay on the old one, budget honestly for the period when both run side by side, and prove the platform with functional and architectural proofs of concept before committing. BearingPoint points insurers toward API decoupling and hybrid delivery. And Datos Insights reports that carriers who planned monolithic suite replacements are reassessing in favor of more composable approaches.

In practice, the combination that tends to work best is **progressive replacement supported by wrap and extend**. APIs keep the legacy core useful while it is still needed, and every new product goes onto a modern foundation from the start. The old system stops growing on the first day and shrinks with every renewal cycle, and nobody has to bet the company on a single weekend.

CHAPTER 6

A Low-Risk Roadmap, One Product at a Time

Modernization done well looks less like a transformation program and more like a series of product launches. Each step is small enough to succeed, visible enough to prove the point, and reversible if something needs to change.

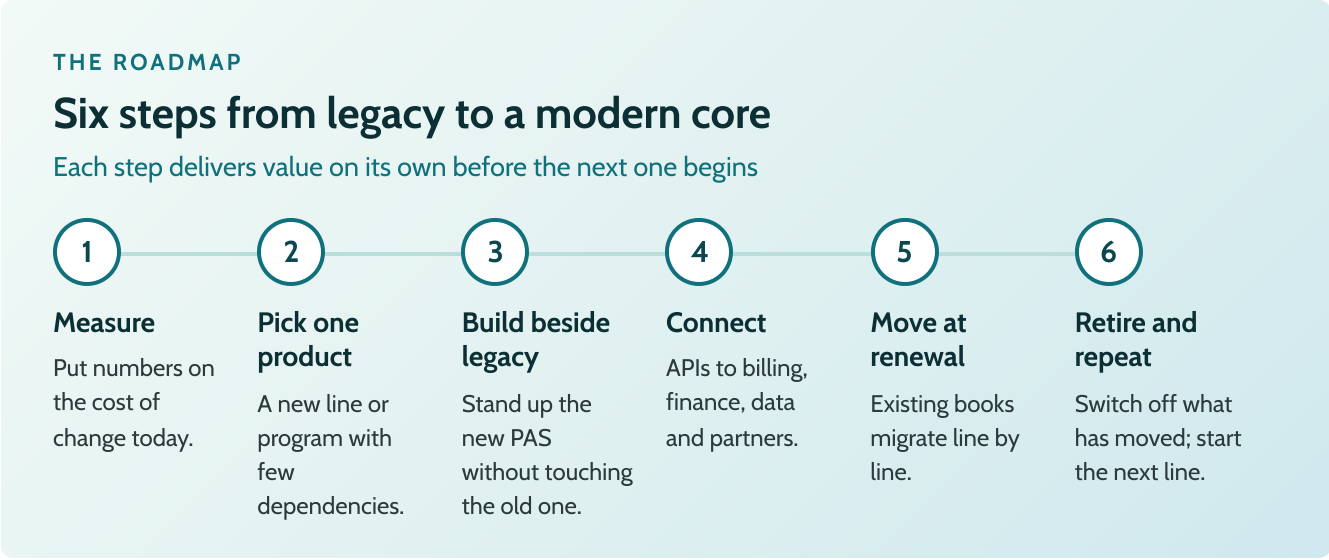

Measure first. Before choosing any technology, put numbers on the cost of the current state: how long a rate change takes, how many hours go into manual work around the core, how long a new product takes from idea to first policy. These numbers become the baseline every later step is judged against.

Pick one product. The best first candidate is a new product, a new program or a small line with few dependencies on the old system. It lets the team learn the new platform on something that matters to the business without risking the book that pays the bills.

Build beside legacy. The new policy administration system goes live alongside the old one, not in place of it. Nothing about the legacy book changes on day one, which removes most of the risk that makes core programs frightening.

Connect through APIs. The new product still needs the general ledger, the data warehouse, payment providers and partner channels. Connecting through APIs rather than batch files keeps data consistent in both directions, and the same connections will serve every product that follows.

Move at renewal. Once the new platform has proven itself, existing books move across one line at a time, usually as policies renew. Each migration is small, testable and easy to schedule around the business.

Retire and repeat. As each line moves, the matching part of the legacy system is switched off. The cost of the old system falls with every cycle, and the team gets faster at each migration.

Three ways to handle the data

Data migration is where programs most often go wrong, so it deserves a decision of its own for each book. **Renewal-based migration** rewrites policies on the new system as they renew, which spreads the work over a year and lets every policy be checked as it moves. **Bulk conversion** suits small, clean books that can move in one step after a full test run. **Run-off** leaves closed or shrinking books on the legacy system until they expire, which is often cheaper than moving them at all. Whichever route a book takes, Roland Berger's advice applies: fix fragmented data before it moves, so its problems don't move with it.

Where AI now makes the journey easier

AI is starting to take the drudgery out of modernization itself. BCG reports that AI-assisted code translation is 40% to 50% faster, and 20% to 25% faster across the whole delivery lifecycle, with code quality 20% to 30% higher. McKinsey's 2026 research on agentic AI in core modernization points to productivity gains of 20% to 50% in discovery and reverse engineering of old systems, and up to 90% in testing and reconciliation. The slowest, most error-prone parts of a migration, understanding what the old system actually does and proving the new one does the same, are exactly the parts AI is getting good at.

CHAPTER 7

Choosing a Modern PAS, and Where Openkoda Fits

A progressive roadmap only works if the platform at its center can carry it. These eight questions separate a policy administration system that will still be easy to change in ten years from one that will become the next legacy.

1 **Can the business configure products, rates and rules without a development ticket?**
Ask to see a rating factor changed and tested live. Every change should be versioned, so policies already bound keep their original terms.

2 **Does it have open, documented APIs that grow with the data model?**
If adding a field means waiting for the vendor to update the API, every integration becomes a project.

3 **Can the data model hold a risk nobody has written yet?**
Specialty, parametric and embedded products rarely fit a fixed schema. Test it with a product the vendor hasn't seen.

4 **Who owns the code, the configuration and the data?**
Ownership and a choice of where the system runs, whether managed cloud, private cloud or on premises, are what keep a platform from becoming the next lock-in.

5 **Can it run beside the legacy system for as long as needed?**
Progressive replacement depends on coexistence: several products, brands or books on one installation, and clean connections to the old core.

6 **Are audit, security and roles built in rather than bolted on?**
Every change should be logged automatically with who, what and when, and access should be controllable down to the field.

7 **Is it ready for AI you can inspect and control?**
Clean, structured data and AI that proposes changes for people to approve are worth more than a chatbot on top of a closed system.

8 **How fast is the first product live, and is the price predictable?**
Ask for a reference that confirms time to first product. Pricing tied to premium or seats makes growth more expensive every year.

One system that brings these qualities together

Few platforms pass all eight tests. Some are easy to configure but closed to integration; others are open but leave the insurance logic for you to build. Openkoda was designed to combine them in one system: a **highly customizable policy administration system** that is quick to stand up beside a legacy core, open to everything around it, and flexible enough to take on each new line as it moves across. Two ideas shape how it is built. **Adaptability** means the system bends to the way your business already works. **Freedom** means that once it does, the code, the data and the choice of where it runs stay with you.

It covers the full P&C lifecycle in one connected set of modules:

- **Product Builder and AI Product Builder.** Define risk types, attributes, coverages and limits on the product itself, or turn a plain-language specification into a reviewable set of changes that nothing applies until a person approves it.
- **Rating Engine.** Rating tables an underwriter can read, formulas shown in plain English, taxes and fees, and a test panel that prices a real risk before anything is saved. Every rate change is versioned.
- **Underwriting rules and Underwriter Workbench.** Decline, referral and auto-approval rules per product, and a workbench where referred cases arrive with the reason attached and the full rating breakdown visible.
- **Policy administration.** Quote, bind, endorse, renew and cancel, with a Policy 360 page that shows coverages, instalments, claims, documents and emails for a policy in one place.
- **Claims processing.** FNOL through settlement, with workflows as simple or as automated as each product needs, including AI-assisted triage, all on the same record as the policy.
- **Billing and payments.** Invoicing, instalment plans, dunning and card payments, recorded and visible on the policy.
- **Forms and distribution.** Quote and claim forms generated from the product definition and embeddable in any website or partner app with one snippet.
- **Workflow Engine and Document Generator.** Visual lifecycle automation that goes live without a release, and versioned Word templates that produce policy documents and branded emails automatically.
- **Reporting, Reporting AI and bordereaux.** Drag-and-drop dashboards, plain-English questions answered with SQL inside your own environment, and premium and claims bordereaux on each program's schedule.
- **Security and control.** Roles down to the record and field, a full audit trail written automatically, and several products, brands or programs on one installation with their data kept separate.

A FOUNDATION BESIDE LEGACY

How Openkoda sits in an insurer's stack

New products go live on Openkoda while the legacy core keeps running what hasn't moved yet

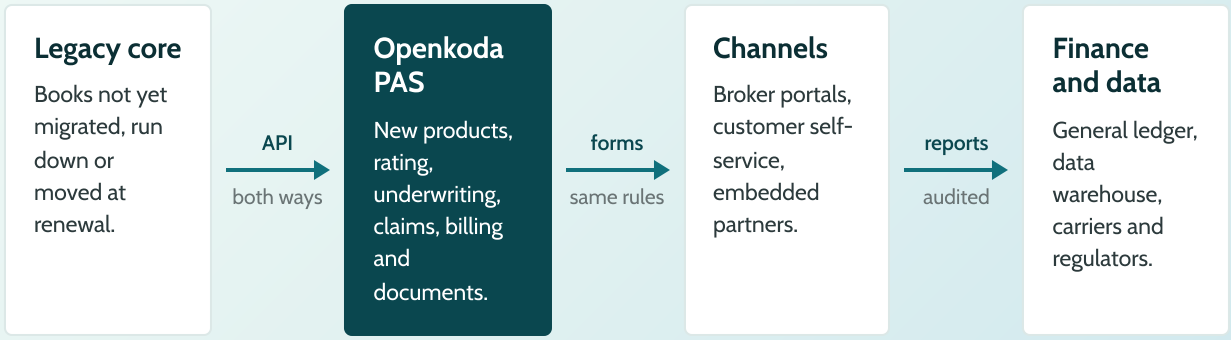

Simplified. Each connection uses Openkoda's REST API or integrations written in standard code.

Integrates with any system around it

Openkoda's REST API is generated automatically and **evolves with the data model**: add an attribute to a product, and it is available through the API without waiting for a vendor release. Custom integrations are written in standard, open code or configured in the browser, so connecting a legacy core, a general ledger, a payment provider, a rating service or a data warehouse doesn't depend on proprietary middleware. Quote and claim forms embed into any website or partner app with a single snippet, and a submission arriving through the API goes through exactly the same validation as one typed into a form. For AI tools, an MCP server exposes the system to LLM clients under the same roles and audit trail as the web application.

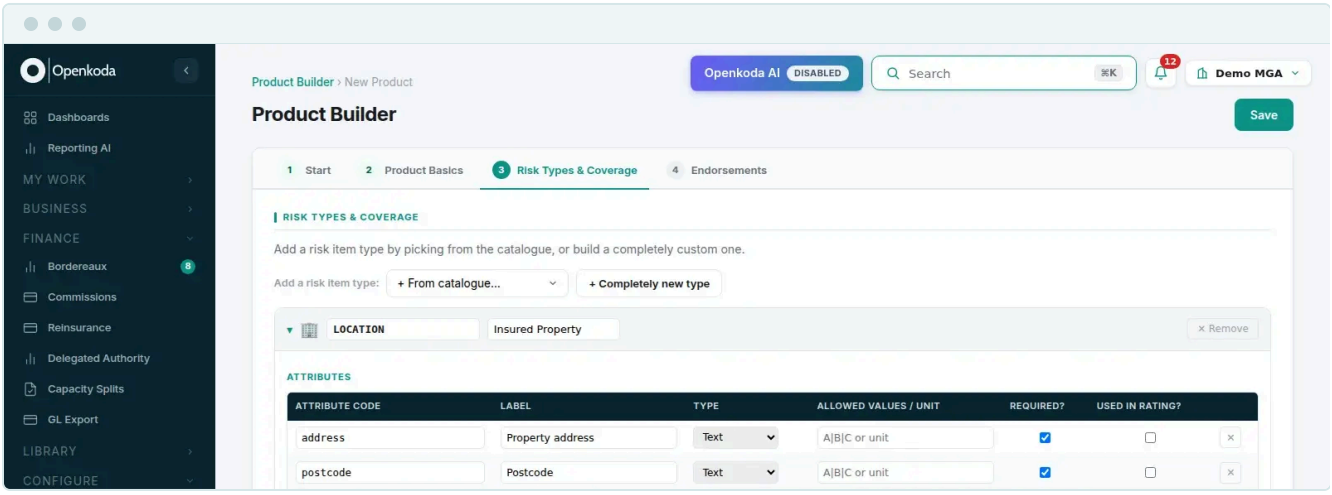

A data model without a ceiling. Risk types, attributes and allowed values are configured on the product, and each new field is immediately available to rating, forms and the API.

Future-proof by design

Openkoda is built on a standard, widely used technology stack: Java and Spring Boot, JavaScript, and PostgreSQL. Any enterprise developer can extend it from day one, and there is no proprietary language to learn or leave behind. It runs in Openkoda's managed cloud, in your own private cloud or on premises, and you can change your mind later. The data model has no fixed ceiling: the same product builder has been used to stand up cover for AI agents as an insured risk, with attributes such as model version and autonomy level, a category that didn't exist a few years ago. Products and rates are versioned, so policies bound last year keep pricing under last year's rules.

> Openkoda was born from a simple idea: insurers should be able to launch and evolve digital products fast, on their own terms, without vendor lock-in.
>
> **Michał Głomba, CRO & Co-Founder, Openkoda**

A foundation for replacement, one product at a time

Because Openkoda can run beside a legacy core, the first product can go live without a big-bang cutover. Several products, brands or programs can share one installation, each with its own people, rules and branding, with data genuinely isolated between them. As books move across at renewal, they land on the same foundation, the same APIs and the same audit trail, and the underwriting and claims teams work in one place instead of two.

Openkoda's own implementation team works with you to shape the system around your products and processes, typically taking a tailored PAS from kickoff to live in **about three weeks**. Pricing is a flat annual rate, with no per-seat fees and no percentage of premium, so the cost of the platform doesn't rise as your book grows.

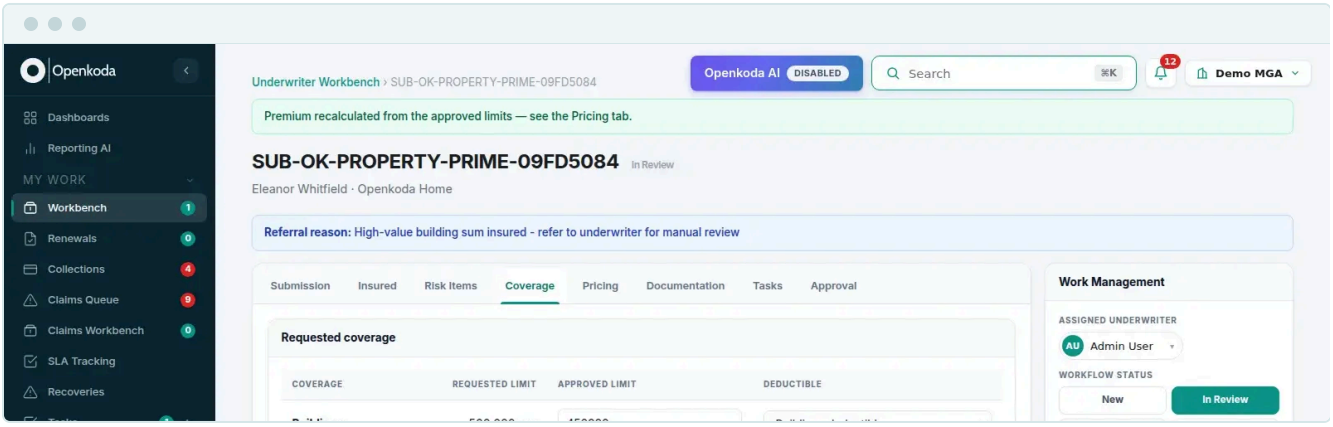

Underwriter Workbench. A referred case with its referral reason, approved limits and a one-click premium recalculation, configured per product.

CONCLUSION

Modernization Doesn’t Have to Hurt

The industry’s fear of core modernization was earned, but it was earned on big-bang programs that tried to change everything at once. Modernizing one product at a time, beside the legacy system rather than in place of it, removes most of what made those programs frightening. Each step is small enough to finish, and each one pays for part of the next.

And the step is worth taking. McKinsey found that insurers on a modern core can run with up to 40% lower IT costs and 40% higher operational productivity. Capgemini’s 2026 research shows the 10% of P&C insurers with advanced AI maturity achieving 21% higher revenue growth and 51% greater share price gains over three years than the rest, and that kind of maturity starts with clean data on a system that is easy to change. Customers notice too: the gap between a 4% and a 52% likelihood to switch is largely a gap in digital experience.

The good years won’t last forever, and neither will the people who understand the old systems. The insurers that use this window to put a modern foundation under their next product will be the ones that can still change quickly when the market turns.

Openkoda

TAKE THE NEXT STEP

Launch your next insurance product in weeks, not years.

Book a personalized demo with Openkoda's implementation team. We'll walk through your own products, rating and underwriting rules, and show you what it takes to run them on a policy administration system you fully own.

- ✓ Live in about three weeks
- ✓ Flat annual rate
- ✓ Your code, your data

openkoda.com/demo

About Openkoda. Openkoda is a highly customizable policy administration system for insurers, MGAs and insurtechs, covering the full lifecycle from quote to claim. It is configured around your products and handed over as yours: your code, your data and your choice of cloud or on-premises hosting, delivered with Openkoda's own implementation team on a flat annual rate.