Openkoda

Choosing the Right Claims Processing System for Your Insurance Business

- ✓ Vendor-Neutral Buyer's Framework
- ✓ 10-Point Vendor Checklist
- ✓ Backed by 2025 and 2026 Research

A mini-ebook by Openkoda

Contents

1. Claims Is the Moment of Truth 3
2. What a Claims System Really Does Today 4
3. Why Claims Projects Disappoint 6
4. Specialty Claims Don't Look Like Auto 7
5. AI in Claims: What Actually Works 9
6. Build, Buy, Extend or Use Your PAS 10
7. A 10-Point Checklist for Vendors 12
8. Where Openkoda Fits 13
9. Conclusion: The Year-Three Test for Claims 16

WHO THIS EBOOK IS FOR

Claims leaders, COOs and buying committees at insurers, MGAs, TPAs and insurtechs who are selecting or replacing a claims system, and everyone who will have to work in it afterwards.

What you will take away: what a claims system really has to do, why claims projects disappoint, how specialty lines and AI change the requirements, four routes to a new system and a 10-point checklist to put to every vendor.

Authors



Arkadiusz Drysch

CEO & Co-Founder, Openkoda



Michał Głomba

CRO & Co-Founder, Openkoda



Arkadiusz Krysik

Marketing Manager at Openkoda

This ebook is distributed under the Creative Commons BY-ND 4.0 License. The license allows for redistribution, as long as the ebook is passed along unchanged and in whole, with credit to its authors.

CHAPTER 1

Claims Is the Moment of Truth

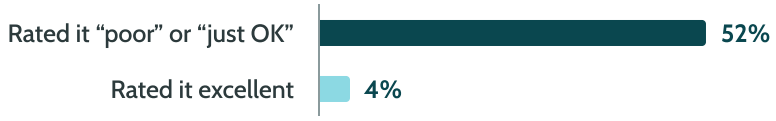
For most customers, a policy is a promise they never test. They pay the premium, receive the documents and forget about it. The claim is the day the promise gets tested, and it is the only day many of them will ever judge you on.

That is why claims is where retention is won or lost. J.D. Power's 2025 U.S. Claims Digital Experience Study asked auto and home customers how their claim went online. Among those who rated the digital experience "poor" or "just OK," **52% said they were likely to leave or not renew**. Among those who rated it excellent, just 4% were at risk of switching. Few levers in an insurance business move retention by that much.

FIGURE 1

A poor digital claim drives customers away

US auto and home claimants at risk of leaving, by how they rated the digital claims experience, 2025



Source: J.D. Power, 2025 U.S. Claims Digital Experience Study (5,958 evaluations), via Carrier Management

The same study found something just as telling: insurers gave customers adequate digital updates on their claim only 22% of the time, even though updates are among the biggest drivers of satisfaction. Customers are rarely angry that a claim takes work. They are angry when they can't see what is happening.

And it still takes a while. In J.D. Power's 2026 property claims study, the average homeowner waited **40.7 days for a final payment** and almost 30 days for repairs to be completed. Both figures improved on the year before, which says less about how fast claims have become than about how slow they were.

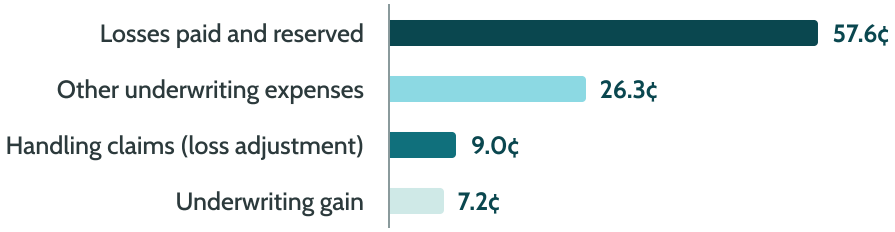
Where the money goes

Claims is also where most of the money goes. In 2025 the US property and casualty industry earned $959 billion in premium. According to the NAIC's year-end figures, $552 billion of it went straight back out as losses, and a further **$86 billion was spent simply on handling claims**: adjusters, investigators, legal fees and the systems that support them. Every other underwriting expense put together came to $252 billion.

FIGURE 2

Where a premium dollar went in 2025

US property and casualty industry, share of net premiums earned



Source: NAIC, P&C Industry Snapshot for the period ended December 31, 2025. Loss adjustment expenses $86.0B on $958.7B earned.

Read those numbers side by side and the case for taking a claims system seriously makes itself. Two-thirds of every premium dollar flows through claims. A few points of leakage (a reserve set too low, a recovery never pursued, a duplicate payment, a fraudulent claim paid in full) cost more than most technology budgets. The underwriting result for the whole year was a little over seven cents on the dollar, which is not much room for a claims operation that runs on rekeying and email.

Then there are the people. The US Bureau of Labor Statistics expects the number of claims adjusters, appraisers, examiners and investigators to **shrink by about 6% between 2025 and 2035**, even as the number of claims does not. The work will have to be done by fewer people, which means the system they work in decides how much each of them can handle.

So a claims system is not back-office plumbing. It sets your loss ratio, your expense ratio and your renewal rate at the same time. The rest of this ebook is a framework for choosing one with that in mind.

CHAPTER 2

What a Claims System Really Does Today

Ask a vendor what a claims system does and you will hear the lifecycle: first notice of loss, investigation, settlement, close. That is accurate and incomplete. A modern claims system runs a chain of decisions, each with money and a regulator attached, and then hands the results to half the business.

THE CLAIMS LIFECYCLE

Seven jobs every claims system has to do well

1. **FNOL**
   Capture the loss from any channel, check cover against the policy, start a complete file.
2. **Triage and assignment**
   Score severity and complexity, fast-track the simple, route the rest to the right handler.
3. **Reserves**
   Set and revise case reserves per coverage line, with a history of every movement.
4. **Payments**
   Pay within each person's authority, escalate the rest, check every payment against the reserve.
5. **Recoveries**
   Spot subrogation and salvage early and track them through to receipt.
6. **Fraud**
   Flag suspicious patterns and send them to investigation, not to payment.
7. **Documents and reporting**
   Letters, notes, evidence and an audit trail, plus the numbers management needs.

Most RFPs stop there. The more revealing question is what happens to claims data once it exists, because a claim is not only a cost to be settled. It is the most honest feedback the business ever gets about its own products.

Pricing needs claims experience by segment, cover and peril to know whether rates are adequate. **Underwriting** needs it to see which risks behave differently from what the guidelines assumed. **Reinsurance** needs incurred figures and large-loss notifications on the treaty's terms and timetable. And for an MGA or a coverholder, **the carrier needs a claims bordereau** every month, in its format, reconciled to the penny. In many organizations each of those feeds is a spreadsheet someone rebuilds by hand, which means each one is late, slightly different from the others and easy to dispute.

That gives a second job description for a claims system, less glamorous than the first: be the single, trusted record of every claim, so that everyone who depends on claims data reads the same numbers.

Configuration, not customization

Every claims vendor will say yes to the seven jobs. What separates them is how much work it takes to change the way each job is done, and who is allowed to do it. Claims rules change constantly: a new authority limit for a senior handler, a new routing rule for water damage after a storm, a new document checklist for a new product, a new status a regulator wants to see. If those changes are

configuration, a claims lead makes them in an afternoon and they survive the next upgrade. If they are **customization**, each one is a development project with a release date, and the claims operation runs on last quarter's rules until it ships.

Keep that distinction in mind. It explains most of what goes wrong in the next chapter.

CHAPTER 3

Why Claims Projects Disappoint

Claims replacements rarely fail outright. More often they go live, work, and quietly fall short of the business case. Handlers keep a spreadsheet on the side. The bordereaux team still rebuilds the numbers by hand. The fast track handles fewer claims than the demo promised. Across core insurance systems in general, Roland Berger's 2026 study found that around half of replacement programs miss their time or budget targets. In claims the causes are remarkably consistent, and there are five of them.

Claims is separated from policy data. When the claims system holds a copy of the policy rather than the policy itself, coverage checks run against data that may be a day, a week or an endorsement out of date. Handlers learn to open both systems and compare. Every mismatch becomes a manual check, and every manual check is a delay the customer feels.



Rigid workflows. Many claims systems encode one way of handling a claim, usually the vendor's idea of a motor or property claim. Everything that doesn't fit gets forced through the same steps, or handled outside the system. The tell is a handler who says "I just skip that screen."

Every rule change needs a release. This is the most expensive pattern and the least visible in a demo. A reserving threshold, an approval limit or a new routing rule should be a setting. When it is code instead, small changes queue behind big ones, and the claims operation adapts to the release calendar rather than to the claims in front of it. West Monroe's 2025 survey of US insurers found that 46% need nine to sixteen weeks for even a minor update to their core systems. In claims, nine weeks is the whole life of a typical file.

Poor integrations. A claim touches payments, documents, repairers, assessors, medical providers, fraud tools, finance and reinsurance. A system that can't exchange data with them cleanly, through an API, turns handlers into couriers who copy information from one screen to another.

AI bolted on. The newest pattern. A separate AI tool reads documents or drafts letters, but sits outside the workflow, can't see the policy and leaves no trace in the claim file. It demos well and changes little, because the handler still has to carry its output back into the system by hand. Chapter 5 looks at what works instead.

WHAT THIS MEANS FOR BUYERS

Ask to see a claim checked against the live policy, a routing rule changed without a release, and an AI suggestion recorded in the claim file. Those three demos tell you more than a feature list.

CHAPTER 4

Specialty Claims Don't Look Like Auto

Most claims software was designed around high-volume personal lines: a car accident or a burst pipe, a repair estimate, a payment. That model is fine for the business it was built for. The trouble is that the fastest-growing parts of insurance look nothing like it. US surplus lines premium, the home of hard-to-place and specialty risk, reached a record $143.3 billion in 2025 according to AM Best, and MGAs, embedded programs and new risks such as cyber and AI liability are all growing faster than the market around them.

Each of these lines bends the claims process in a different direction.

Parametric. There is no loss adjustment in the usual sense. The policy pays a fixed amount when an agreed index crosses a threshold: wind speed at a location, rainfall over a season, an earthquake's magnitude, a flight's delay. The claims system has to ingest the trigger data, match it to every affected policy and pay quickly, often without the policyholder filing anything. The hard parts are data feeds, bulk processing and basis-risk disputes, not investigation.

Travel. Many small claims, a few very large ones and a lot of urgency. A medical emergency abroad involves assistance providers, foreign invoices in other currencies and guarantees of payment issued before the claim is even complete. Group programs add another layer: which organization, which traveler, which upgrade, on which dates.

Aviation. Low frequency, high severity, several coverage lines on one event (hull, liability, passenger injury) and often several parties on the risk. Reserves move in large steps, reinsurers want to know early, and the file can stay open for years.

Embedded. The claim often starts inside someone else's product: a booking app, a checkout, a device warranty flow. FNOL arrives through an API, the claims are small and numerous, and the partner wants status updates in its own interface. Straight-through processing isn't a bonus here; it's the business model.

New risks. Cyber and AI liability claims bring coverage lines nobody had five years ago: a chatbot that invents a refund policy, a model that leaks personal data, an agent that infringes a copyright. The evidence is a transcript and a tool-call log, not a photo of a dented bumper.

The comparison table sums up how far apart these lines really are.

TABLE 1

How specialty claims differ

Line	How claims start	What decides them	What the system must handle
Personal auto or home	Customer call, app or broker	Damage estimate, coverage check	High volume, repair networks, fast track
Parametric	Index crosses a threshold	Trigger data, not adjustment	Data feeds, bulk payment, no FNOL
Travel	Assistance call from abroad	Medical and trip evidence	Foreign currency, guarantees, group cover
Aviation	Broker or operator notice	Surveyors, experts, counsel	Multiple lines per event, large reserves
Embedded	Partner API	Simple rules, small amounts	Straight-through processing at scale
Cyber and AI liability	Incident notice	Logs, forensics, legal review	New coverage lines and evidence types

The lesson for buyers isn't that you need a different system for every line. It's that the system has to let you define **claim types** (their data fields, statuses, documents, routing and authority) per product, and add new ones as the book changes. A claims system that can only handle the claim types it shipped with will force every specialty line into a shape that doesn't fit it, and the gap will be filled with spreadsheets.

CHAPTER 5

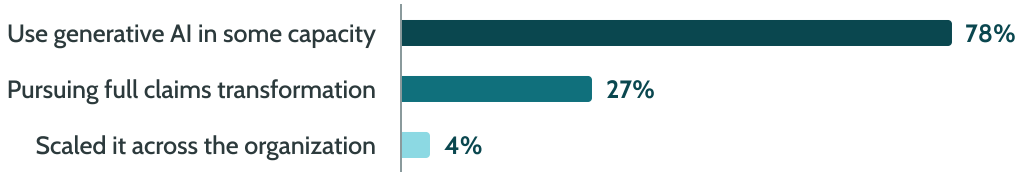
AI in Claims: What Actually Works

No area of insurance gets more AI promises than claims, and few have a wider gap between promise and practice. Bain & Company's 2025 Claims Maturity Assessment of 81 P&C insurers found that 78% were already using generative AI in some capacity. Only 4% had scaled it across the organization. The pattern isn't unique to insurance: MIT's NANDA initiative reported in 2025 that only about 5% of enterprise generative AI pilots achieve rapid revenue gains.

FIGURE 3

Generative AI in claims: widely tried, rarely scaled

Share of 81 P&C insurers worldwide, 2025



Source: Bain & Company, 2025 Claims Maturity Assessment, via Insurance Business

The insurers getting real results have something in common: they give AI **narrow jobs inside the workflow**, not a broad mandate outside it. Allianz's Project Nemo in Australia is a good example. It handles one small, well-defined kind of claim: food spoiled during a storm-related power outage, under AUD 500. Seven task-specific agents each do one job (checking coverage, confirming the weather event, looking for fraud, calculating the payout, summarizing the case), and then **a human claims professional makes the payout decision**. Claims that used to take several days are settled in a day or even hours, and Allianz reports an 80% reduction in processing and settlement time. At Travelers, more than half of all claims are now eligible for straight-through processing, and customers choose it about two-thirds of the time when it is offered.

Neither is a story about AI replacing claims handlers. Both are stories about AI clearing the simple, repetitive work so that people can spend their time on the claims that need judgment.

Keep a person on what matters

Customers are clear about where they want that line drawn. In a KPMG survey of 2,000 UK adults in March 2026, **90% said human interaction is important in claims handling**, and 64% said claims should be handled primarily by people. Fifty-five percent were not confident that fully automated handling would produce fair and accurate outcomes.

Regulators are moving the same way. Nearly half of US states had adopted the NAIC's model bulletin on insurers' use of AI by early 2025, which expects insurers to govern, document and test the AI systems behind decisions that affect consumers. In Europe, the AI Act already treats AI used for risk assessment and pricing in life and health insurance as high risk. Whatever the exact rules in your market, the direction is consistent: you will need to show what the AI did, on which data, and who approved the outcome.

DRAWING THE LINE

Where AI earns its place in claims, and where a handler decides

LET AI DO THE GROUNDWORK

- **Read and sort** documents, photos and emails at intake
- **Pre-fill** the FNOL file and spot missing information
- **Score** severity and fraud risk to guide triage
- **Draft** letters and requests in your house voice
- **Summarize** long files for the next handler

KEEP A PERSON IN CHARGE

- **Coverage decisions** and every denial
- **Reserves** above set thresholds
- **Payments** outside fast-track rules
- **Fraud referrals** before any action
- **Anything sent** to a claimant in a dispute

For a claims system, this translates into four practical requirements. AI must work **on the claim file**, with access to the live policy and claim history, not on a copy pasted into a chat window. Its output must be a **suggestion a handler reviews**, whether that is a draft letter, a triage score or a reserve recommendation. The review and the final decision must be **recorded in the audit trail**. And the business must be able to decide, claim type by claim type, **where the line between automatic and human sits**, and move it as confidence grows.

CHAPTER 6

Build, Buy, Extend or Use Your PAS

Before comparing vendors, decide which route you are on. There are four, and they lead to very different places three years after go-live.

Build your own. A claims system shaped exactly to your process, owned outright. You also take on everything that isn't claims: permissions, audit, documents, payments, integrations and the people to maintain all of it. It makes sense for a large carrier with a permanent engineering team and a claims process that is genuinely a competitive advantage. For most others, it is a long road to a system that does what others already do.

Buy a standalone claims suite. Deep claims functionality, a long reference list and mature integrations with repair networks and assessors. The trade-off is the boundary with the rest of your business. A standalone suite holds its own copy of policy data, so the coverage check, the bordereau and the finance reconciliation all depend on integrations that someone has to build and keep in sync. Ask how rule changes are made and who makes them.

Extend what you have. Add a workflow tool, a portal or an AI service on top of the existing claims system. It is the cheapest and fastest option for a specific pain point, and often the right first step. But each extension adds another place where claims data lives, and the underlying limits (rigid workflows, release-driven change) stay where they were.

Use the claims module in your PAS. Claims, policy and billing share one data model, so cover is checked against the live policy, payments and recoveries land in the same ledger and bordereaux come from the same records. The question to ask is depth: is the claims module a genuine claims system, with reserves, authority, recoveries, fraud and a handler workbench, or a claims register attached to a policy system?

The comparison table summarizes the trade-offs.

TABLE 2

Four routes to a claims system, compared

Dimension	Build	Standalone suite	Extend existing	Claims module in PAS
Time to first live claims	Many months	Often a year or more	Weeks	Weeks to months
Cover checked against live policy	If you build it	Via integration	Unchanged	Yes, same data
Rule changes without a release	If you build it	Varies by vendor	Unchanged	Varies by vendor
Bordereaux and finance from one record	If you build it	Via integration	No	Yes
Depth of claims functionality	What you build	Deep	As before	Check carefully
Ongoing engineering burden	High	Vendor and integrator	Grows over time	Lower

No route is right for everyone. A carrier with a stable, high-volume personal lines book and a strong integration team may be well served by a standalone suite. An insurer with one urgent pain point should probably extend first. But for organizations whose products change often (MGAs, specialty carriers, insurtechs and anyone launching new lines), the boundary between policy and claims is where most of the cost and delay hides, and removing it is worth a lot. Whichever route you lean toward, the checklist in the next chapter applies to all of them.

CHAPTER 7

A 10-Point Checklist for Vendors

Every claims demo shows a smooth claim moving from FNOL to payment. The ten questions below are designed to find out what happens to the claims that aren't smooth, and to the rules that will need to change after go-live. They fall into three groups: the data the system works on, how easily it changes, and how much control you keep. Send them in writing before the demo, then use the demo to check the answers.



Data

1 **Is cover checked against the live policy at FNOL?**
Ask for a claim on a policy endorsed yesterday. If the check runs on a copy that is synchronized overnight, handlers will keep checking by hand.

2 **Which channels can open a claim, and do they all produce the same complete file?**
Web, email, phone, broker portal and API should all land in one validated intake. A claim that starts incomplete starts with a chase.

3 **Can claims data feed bordereaux, reinsurance and pricing without a spreadsheet in between?**
Ask to see last month's claims bordereau produced from the system, in the carrier's format. If someone rebuilds it by hand, it will be late and disputed.



Change

4 **Can a claims lead change routing rules, reserving thresholds and approval limits without a release?**
Ask them to change one live, during the demo. If it needs a developer, your claims rules will always be a quarter behind your claims.

5 **Can you add a new claim type, with its own fields, statuses and documents?**
Bring a specialty line from chapter 4 and ask how long it would take. This is where most systems show their limits.

6 **How does work reach handlers, and can the queue be reordered?**
One queue ordered by SLA, severity or exposure, with automated assignment, absorbs a catastrophe surge. Separate inboxes per handler don't.

Control

7 **Are reserves and payments governed by authority limits, with a full history?**
Every reserve movement and payment should show who made it, under what authority, and against which coverage line.

8 **Are recoveries and fraud built in or bolted on?**
Subrogation, salvage and fraud flags should be part of the claim file, from FNOL onward, not a separate tool someone remembers to check.

9 **Where does AI act, and who approves what it produces?**
Ask to see an AI-drafted letter or triage score, the handler's review, and the record of both in the audit trail.

10 **What will it cost at three times your volume, and where can it run?**
Pricing per user or per claim grows with your success. Ask for the price at triple your volume, and whether the system can run in a managed cloud and on your own infrastructure if data residency requires it.

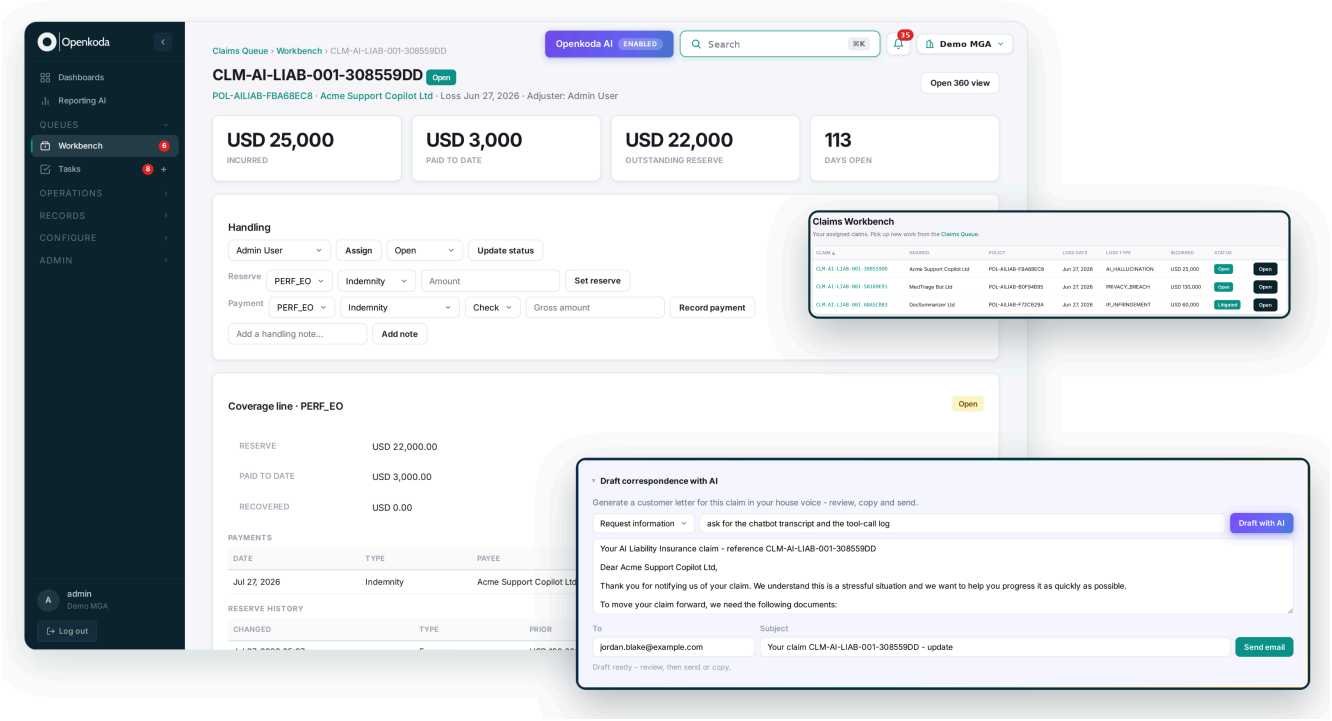
None of these questions is exotic. The value is in asking all of them, in writing, of every vendor, and comparing the answers side by side instead of comparing demos.

CHAPTER 8

Where Openkoda Fits

So how does Openkoda answer the ten questions? It deserves exactly the same scrutiny as everyone else on your list, so what follows shows what sits behind each answer.

Openkoda is a highly customizable, AI-enabled policy administration system with a full claims module. It sits on the fourth route from chapter 6: claims shares one data model with policy and billing, so when a claim arrives, **cover is checked against the live policy**, not a copy. Claims can be opened from the web, email, phone, a broker portal or an API, and every channel feeds the same validated intake, so handling starts from a complete file.



Claims workbench. A claim with its incurred, paid and reserve figures, handling controls per coverage line, the handler's queue and an AI-drafted letter awaiting review.

Handlers work from **one queue**, ordered by SLA, severity or exposure. The workbench brings documents, notes, correspondence and every reserve movement into one view. Reserves and payments are tracked per coverage line, payments are checked against the reserve and routed for approval by role or amount, and the full history stays on the file. Fraud flags, subrogation and salvage are built into the claim rather than handled in another tool.

Configuration, all the way down

The point chapter 2 made about configuration is where Openkoda is most different. Claim types, data fields, statuses, routing rules, reserving thresholds, approval limits and the workbench layout are all configuration. Changing them **needs no redeployment**. A new authority limit after a reorganization, a new routing rule for a storm surge or a new claim type for a specialty product is a change a claims lead makes, not a release someone schedules.

> 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

AI follows the rule from chapter 5. It works inside the claim, with the policy and the file in front of it. Settlement and request letters are drafted with AI in your house voice, from the claim itself, and **a handler reviews every draft before it is sent**. Nothing goes to a claimant without a person deciding it should.

For MGAs and coverholders, claims data **feeds bordereaux directly**, produced on each program's schedule from the same records the handlers work on (see openkoda.com/bordereaux-reporting). Openkoda runs in a managed cloud or on premises, on flat, published pricing that doesn't change with the number of users (openkoda.com/pricing).

In practice

Two published case studies show the range. **SkyGuard** used Openkoda to run life and specialty aviation insurance on one platform in about three months, with FNOL, claims registration and tracking across the claim lifecycle, a dedicated portal for claims specialists and integration with external claims systems where needed. **Core Travel Insurance**, a US provider of group travel medical insurance, replaced a legacy system that couldn't support its distribution model and went from implementation to production in eight weeks.

OPENKODA AGAINST THE CHECKLIST

Ten questions, ten answers

1 **Cover checked against the live policy**
Claims shares one data model with policy and billing.

2 **Channels into one complete file**
Web, email, phone, broker portal or API, validated at intake.

3 **Bordereaux without spreadsheets**
Claims data feeds bordereaux directly, on each program's schedule.

4 **Rules changed without a release**
Routing, reserving thresholds and approval limits are configuration.

5 **New claim types**
Claim types, fields, statuses and workbench layout are configurable.

6 **One queue**
Ordered by SLA, severity or exposure.

7 **Authority and history**
Payments routed by role or amount, every reserve movement recorded.

8 **Recoveries and fraud built in**
Fraud flags, subrogation and salvage on the claim itself.

9 **AI with a handler in charge**
AI drafts settlement letters; a handler reviews them before they are sent.

10 **Cost and deployment**
Flat, published pricing; managed cloud or on premises.

Don’t take any of this on trust. Bring your most awkward claim type to a demo and watch it being configured. More detail is at openkoda.com/claims-management-software.

CONCLUSION

The Year-Three Test for Claims

Every claims evaluation eventually becomes a demo comparison: whose workbench looks cleaner, whose FNOL form is smoother, whose AI writes the nicer letter. Those comparisons feel decisive in the room and matter surprisingly little afterwards.

What no demo shows is year three. By then a storm has tested your queue. A new product has brought a claim type nobody planned for. A carrier wants a different bordereau, a regulator wants to know how your AI made a decision, and two senior handlers have retired. The real test of a claims system is whether your own team can absorb all of that by changing settings, or whether each change waits for a release.

So judge every option by year three, not by day one. Ask the ten questions of the system you already like best, not just of the ones you are comparing it with. Watch a rule change live. And remember what is at stake: the one moment when your customers find out what your promise was worth.

BEFORE YOUR NEXT VENDOR MEETING

Three steps that sharpen any claims evaluation

1

Pick your hardest claim type

The specialty line or new product, not the simple motor claim. Score every vendor on it.

2

Send the ten questions first

Get written answers, then use the demo to check them.

3

Change a rule live

An approval limit or a routing rule, during the demo. Time it.

If you’d like to see how Openkoda handles your claims, book a demo at openkoda.com/demo.

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.

© Openkoda · openkoda.com · Distributed under the Creative Commons BY-ND 4.0 License.