

# Ambitious AI, kept compliant

*How insurers can run powerful generative AI on the hardest work in the business while keeping the compliance and audit trail a regulator requires · Axxion Claims Settlement Services · White paper v1.0 · 31 May 2026*

Executive summary	2
1. The opportunity: the real value is in the ambitious work	3
2. The collision: why ambition and regulation pull apart	4
3. The resolution: put guardrails around the model	5
4. Proof: the seven-stage claims pipeline	6
5. How far guardrails let the model go	7
6. The payoff: auditable by design	7
About Axxion	8



## Executive summary

Generative AI delivers real value in insurance, and its biggest returns sit in the most ambitious work, well beyond the simple tasks. A modern large language model can read and cross-check an entire claim file, find patterns across thousands of adjuster notes and call recordings, summarize long medical and surveyor reports, and draft and translate at scale. These are large, valuable jobs, and the technology is good at them.

The challenge is regulations. A generative model is **probabilistic**: it does not give exactly the same answer twice, and it cannot prove why it produced a given output. Insurance runs on the opposite requirement. Every coverage decision, every payment, and every entry in the record has to be consistent, explainable, and open to inspection years later. The moment a probabilistic model decides coverage or releases money, the audit trail (the record showing how each decision was made) becomes "the model said so," and no regulator accepts that.

Faced with this, most insurers do one of two things: shrink the AI down to trivial tasks, or abandon the pilot. Both leave the value on the table. **There is a third option that keeps the ambition.** Run the powerful model for the analysis, then place guardrails around everything it produces. The guardrails are **deterministic** code: conventional software and rules that follow fixed logic, return the same result every time, and can be inspected and reproduced. That code checks the model's output, scores how confident it is, records it, and makes the decision. Control of decisions, money, and the record never leaves it.

The rule is short: the model proposes, and deterministic code disposes. Axxion runs it across a seven-stage motor-claims pipeline (section 4), and the result is the full reach of generative AI with the three properties regulated claims work requires: compliance, explainability, and containment. The more ambitious the model, the more the guardrails matter, because the guardrails, not the model, hold control.



## 1. The opportunity: the real value is in the ambitious work

The largest returns from generative AI in insurance come from the work conventional software has never been able to touch: making sense of unstructured information at scale. The simple uses, such as pulling a date off a form, are real but small. The value sits higher up.

<b>Ambitious use case</b>	<b>What the large language model does</b>	<b>Example in motor claims</b>
<b>Document intelligence at scale</b>	Reads, structures, and cross-checks whole files: policy wordings, medical reports, surveyor reports, legal correspondence.	Reconcile a repair estimate against the policy terms, the damage photos, and the repair method in one pass.
<b>Text and pattern analytics</b>	Finds themes and signals across large volumes of free text and transcripts.	Surface emerging fraud patterns, leakage drivers, and complaint themes across thousands of adjuster notes and call recordings.
<b>Multi-source reasoning</b>	Compares several documents and data sources at once and flags what does not line up.	Check a new claim against prior claims on the same vehicle, the invoice, and the estimate for inconsistencies.
<b>Portfolio review</b>	Summarizes and triages large claim volumes for review.	Scan closed claims for subrogation and recovery opportunities a manual review would miss.
<b>Language at scale</b>	Drafts and translates across languages.	Produce policyholder and workshop communications in Arabic and English from the same case file.

All of this is cognitive work (LLM): interpretation, analysis, and language. It is exactly where a generative model is strong, and none of it requires the model to make a decision on its own. The mistake many insurers make is to look at the compliance risk and respond by shrinking the AI to the trivial tasks. That keeps the program safe and small. The better course is to run the ambitious work and control it, which is what the rest of this paper describes.



## 2. The collision: why ambition and regulation pull apart

---

*Generative models are powerful because they are probabilistic, and that same property disqualifies them from the regulated core of insurance. The technology that is best at interpretation is the wrong technology for decisions, money, and record.*

---

Two kinds of software sit behind any claims process, and they behave in opposite ways.

Property	Deterministic code (including the insurer's existing models)	Generative LLM
Same input, same output?	Yes, every time	No, the answer varies
Can a past result be reproduced?	Yes	Not reliably
Can it show why it produced the answer?	Yes, traced to the rule or the model factor	No faithful trace; any explanation it gives is itself generated text
Best used for	Decisions, calculations, the record	Reading, interpreting, analyzing, and drafting

Compliance and audit in financial services rest on the deterministic column. A regulator expects every decision to be consistent, explained against a rule that was in force at the time, and reproducible on inspection. A generative model offers none of these for a decision. It can write a fluent account of its reasoning, but that account is generated after the fact and is not a record of how the output was formed. Put the model in charge of a coverage call or a payment, and the audit trail reduces to an assertion that the answer was correct.

### Why the insurer's existing models stay in place

The models an insurer already runs for pricing, fraud scoring, and severity are deterministic at the point of prediction, and they can be explained. That is what makes them defensible to an actuary and a regulator.

Model type	How it is explained	Example
Intrinsically interpretable	The model exposes the weight of each factor directly	Linear model, scorecard, small decision tree
Explained after the fact	A separate, established technique attributes the prediction to its input factors	Gradient-boosted trees with feature attribution
Generative LLM	No reliable attribution; its explanation is generated text, not a record of how the output was formed	A chat-style model drafting a summary

### Replacing those models with a generative one trades an explainable system for an unexplainable one.

The ambitious model belongs on top of them, reading and analyzing, while the interpretable models keep doing the predicting. WTW describes building actuarial tooling the same way: generative AI on top, a deterministic core underneath that a human can open and drill into. (*Pardeep Bassi, WTW Talking Technology.*)



### 3. The resolution: put guardrails around the model

---

*The model can be as ambitious as the insurer wants, if deterministic code guards the way in and the way out and keeps every decision for itself. The model proposes; the guardrails dispose.*

---

A guardrail wraps each model call between two deterministic steps. The workflow engine, built from deterministic code, owns the process from start to finish and calls the model only as a contained step inside it.

Guardrail in (deterministic)	The model call (probabilistic)	Guardrail out (deterministic)
A specific question with an allowed answer format. The input is shaped and constrained before the model sees it.	The model is asked questions, however ambitious, and returns a structured answer. It holds no power to act.	Format check, confidence score, and an audit-log entry. The code decides what to accept, reject, or send to a person.

Three properties make the guardrail hold in a regulated setting.

- **Deterministic code owns the process.** It is the only thing that creates a claim, changes its status, releases a payment, talks to the regulator, and writes the audit log. The model is called for an answer and then set aside.
- **Every model output is scored and tagged.** Each answer is checked against the expected format, scored for confidence, and recorded with the model version, the time, and a flag marking it as AI-generated. A low-confidence answer, or one with no clean match, goes to a person.
- **The audit log is the record.** For every decision it stores the rule version in force, the input received, the model version, the confidence score, and the guardrail's decision.

Building one of these automations is mostly an exercise in defining the guardrail, not the model: the input it allows, the single question it asks, the answer format and confidence threshold it accepts, and where it sends anything that fails. The people who build these models describe the same priority.

*"Trust is mostly a governance and engineering problem rather than a model-intelligence problem."*  
André Balleyguier, Anthropic · WTW Talking Technology



## 4. Proof: the seven-stage claims pipeline

[Axxion Claims Settlement Services L.L.C.](#) runs the guardrail across its full motor-claims pipeline. At every stage the model proposes and deterministic code decides, commits, and records. The model does ambitious reading and analysis throughout; it approves nothing and moves no money anywhere.

Stage	Cognitive work (LLM): what the model proposes	Control work (deterministic code): what the guardrail decides and records
<b>1. FNOL &amp; registration</b>	Parse free-text incident into fields, clean voice transcripts, handle Arabic-English code-switching.	Policy lookup and coverage validation, format gate on every field, workflow engine creates the claim.
<b>2. Triage &amp; assignment</b>	Grade severity from photos, identify make, model, and year from images.	Rule-based routing on capacity, SLA, and certification; severity is one input of many; rule version logged at the decision.
<b>3. Estimation</b>	Read estimate PDFs into line items, identify damaged parts from photos, flag variance against the repair method.	Parts priced by catalog lookup; an unmatched part routes to a surveyor and is never a guess; low confidence routes to review.
<b>4. Authorization</b>	Draft the authorization narrative, summarize photo evidence, draft supplement justifications.	Hard-coded authority matrix, approval as a deterministic state machine, total-loss threshold calculated by rule.
<b>5. Repair &amp; QC</b>	Check milestone photos, classify the pre-delivery inspection, flag anomalies in paint, emblems, and filler.	Photo checklist enforced by the workflow, recalibration certificate checked against the issuer, failed grades routed to human QC.
<b>6. Settlement</b>	Read invoices into line items, detect invoice-versus-estimate gaps, draft settlement notices.	Three-way match releases payment; VAT, ledger posting, and bank rails are deterministic; no model touches money.
<b>7. Close &amp; analytics</b>	Draft the claim summary, classify type and root cause, write closure letters by language.	Closure rules on completeness and retention, burning cost from the actuarial engine, KPIs from the single source of truth with no model involved.

Estimation shows the split clearly. The model reads the estimate PDF and flags where it varies from the expected repair method, ambitious analytical work across a long document. A catalog lookup then prices the parts, and a confidence threshold sends low-certainty cases to a surveyor. The analysis is the model's; the decision and the record are the code's.



## 5. How far guardrails let the model go

Stronger guardrails allow more ambition, not less. The same design supports four levels of use, with a named person accountable throughout.

Level	What the model does	What the guardrail holds
<b>1. Automate</b>	Speeds up existing work: extract, translate, structure.	Code validates and stores every output.
<b>2. Predict</b>	Explains and interprets the predictions the deterministic models make.	The interpretable models still make the predictions.
<b>3. Suggest</b>	Proposes an action or recommendation.	A person decides; code records the decision.
<b>4. Act</b>	Takes a bounded action within set limits.	Code holds the limits and the audit log; a named person stays accountable.

The level changes how much the model proposes. It never moves the decision, the money, or the record off the deterministic code. An insurer can raise the level on a given task as its confidence and its guardrails mature, which is how ambition grows safely. (*Maturity model: WTW.*)

## 6. The payoff: auditable by design

The guardrail around the model buys three properties a regulator and an internal auditor can verify.

Property	What it means	What the audit log shows
<b>Compliance</b>	Every decision is made by deterministic code against versioned rules.	The rule version in force, the input it received, and what it produced.
<b>Explainability</b>	Every model output is traceable and marked as AI-generated.	The model version, the input, the confidence score, and the AI-generated flag.
<b>Containment</b>	No model changes a status, releases payment, approves a claim, or sets a KPI.	The guardrail's decision: accept, reject, or escalate.

These properties are operational protection rather than paperwork. A single wrong output from an otherwise reliable model can end a program's credibility, and trust is far harder to recover than to lose. The guardrail keeps one bad answer from ever reaching a decision, a payment, or the regulator record. That is what lets an insurer be ambitious with the model and stay defensible at the same time: the model proposes, and deterministic code and its guardrails dispose.

*This document is provided for general informational purposes and reflects Axxion Claims Settlement Services' operational approach as of 31 May 2026. It does not constitute legal, regulatory, or actuarial advice. Architectural descriptions are illustrative of Axxion's claims-handling design and may differ from any specific insurer implementation. Frameworks and statements attributed to third parties (Anthropic, WTW) are the work of those organizations, cited here for context and not presented as Axxion data; readers should consult the primary sources before relying on them. Regulatory requirements vary by jurisdiction.*



## About Axxion

[Axxion Claims Settlement Services L.L.C.](#) is the UAE's first dedicated motor third-party administrator. From Dubai, Axxion manages the full motor claims lifecycle for insurance partners: first notification of loss, surveying, repair coordination, quality control, recovery, and settlement. The company serves UAE insurers across both large and small-to-medium carriers and is preparing to extend into Saudi Arabia and the wider GCC.

**Compliance by design.** Axxion was built to operate inside a tightening regulatory environment. The Central Bank of the UAE absorbed insurance regulation in 2020 and consolidated the framework under Federal Decree-Law No. 6 of 2025, which brings TPAs and loss adjusters explicitly inside the CBUAE perimeter. Every claim Axxion handles passes through formal compliance gates covering UAE PDPL data protection, policyholder-consent requirements, settlement-authority bands, sanctions screening, and audit-trail completeness. Compliance is not an overlay; it is the operating substrate.

**The Axxion Intelligent Operating System (AIOS).** The Claims OS that Axxion presents to insurer partners runs on the AIOS, a unified operating layer orchestrating a seven-stage claims pipeline across surveying, estimation, repair coordination, quality control, recovery, settlement, and reporting. The AIOS integrates human operators with structured AI-assisted decision points at every stage. Insurers receive cleaner data, faster cycle times, lower per-claim cost, and a complete audit trail, without giving up control over their portfolio.

Axxion is led by Managing Director and Co-Founder [Frederik Bisbjerg](#). Frederik has spent over two decades in international insurance in operating and transformation roles, with a particular focus on motor claims, distribution, and AI-enabled insurance models. He has been a public voice on insurance modernization in the GCC and writes regularly on the operational shifts reshaping motor underwriting. Axxion is the operating expression of his thesis: that disciplined claims operations, built on AI-enabled but human-led decisioning, will define the next decade of competitive advantage in motor insurance.

### More information:

Managing Director and Co-Founder Frederik Bisbjerg | [f@axxion.co](mailto:f@axxion.co) | [www.axxion.co](http://www.axxion.co)