sprout.ai

2026



AI FOR SPECIALTY LINES INSURANCE CLAIMS

Run both high-volume and complex specialty claims on one operating model

OVERVIEW

The same depth of coverage check on every claim, regardless of its damages and complexity

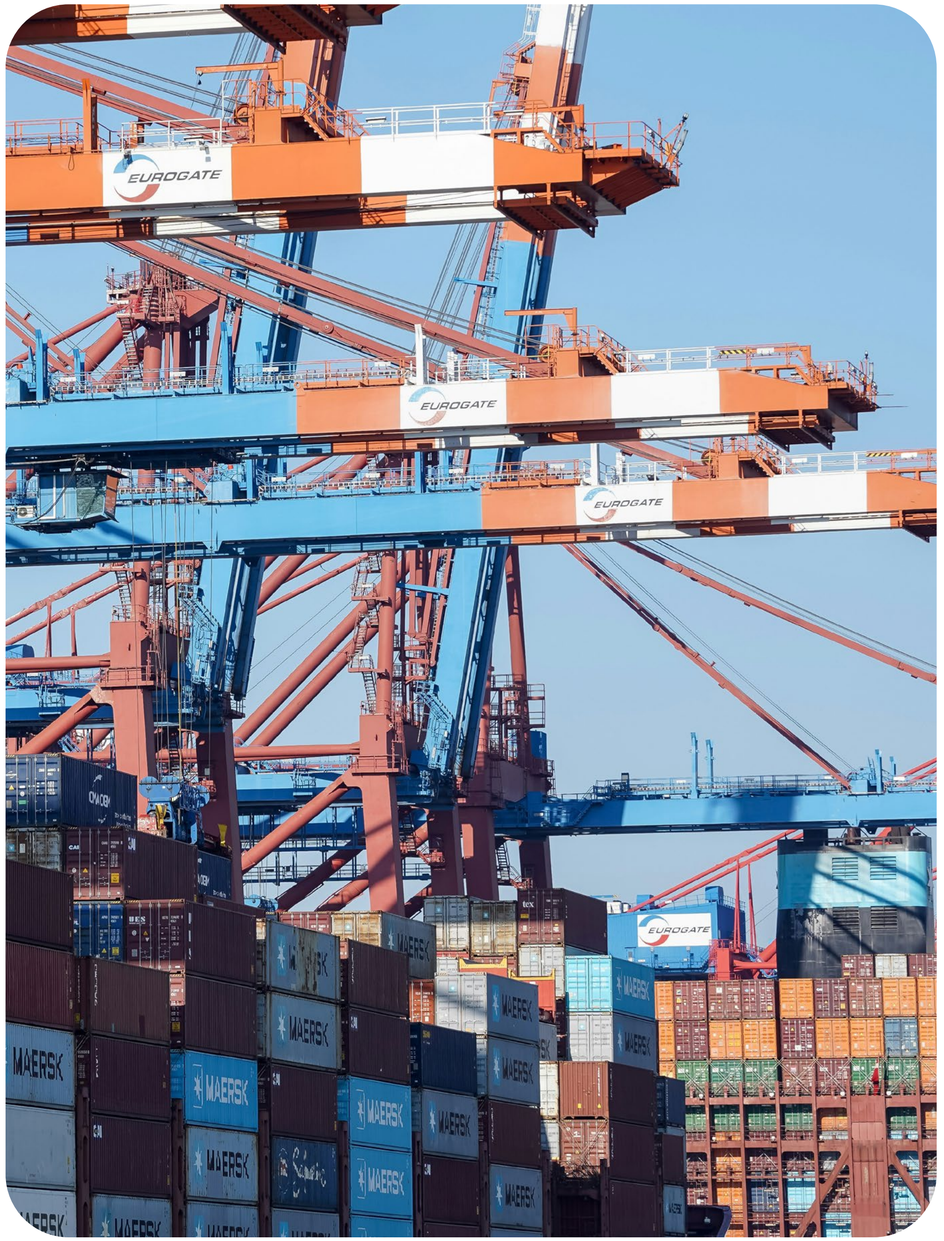
A specialty portfolio can carry two claim populations. One is high in volume and low in value, where the constraint is throughput. The other is low in volume and high in value, where the constraint is how deeply the wording and the evidence get read. Both are decided against bespoke terms, both can move years after the claim is first notified, and the same team is accountable for both.

They are not separate books. A shortage turns into a subrogation or general average file. Cargo follows the same pattern: 65% of cargo damage claims come down to poorly packed, blocked or secured containers, and individual losses can still run into eight figures. A professional indemnity notification turns into a coverage dispute once the pleadings arrive.

Which path a claim takes is settled early, on a limited account of what happened, and the cost of getting it wrong runs in both directions. Specialist time is spent on claims that never needed it, and indemnity paid on claims that were never checked properly.

Sprout.ai reads the policy, the endorsements and the evidence together, identifies which terms govern, and recommends the coverage position with the clause cited. The same check runs on every claim in the book, so a routine file is worked as a routine file and a developing one surfaces as soon as the evidence says so.

Sprout.ai recommends. Your adjuster decides. Every recommendation carries the governing document, the clause applied and the reasoning, so any outcome can be reviewed, overridden, or defended in a dispute.



THE CHALLENGE

One specialty book, two kinds of claim — and they cannot be worked the same way

Specialty claims require detailed review, but reviewing every claim manually limits how efficiently teams can handle volume and complexity.



Two operating modes, one book

High-volume claims need a fast, consistent path through notification, cause classification and coverage. Complex files need depth. Bills of lading, expert reports and loss calculations read together against layered wording. Carriers end up running both, often on separate systems with separate authority levels and separate queues. The specialist capacity the complex files depend on is spent supervising the volume ones.

Volume claims become complex claims

Specialty claims range from high-volume, lower-complexity cases to complex claims requiring specialist review. A straightforward liability notification becomes a coverage dispute once the pleadings arrive. The classification that determines which path the claim takes is made when the least is known about it, and revisiting it later is what turns a small claim into an expensive one. The same pattern holds in cargo: the most costly 1% of marine claims account for at least 30% of total claims value each year. *Source: Allianz Commercial*, “Loss trends - Larger vessels bring bigger losses”.

Documentation arrives from every direction, in every format

Carriers, shippers, surveyors, brokers, insureds and lawyers all submit, in inconsistent formats and multiple languages, across jurisdictions with different legal and regulatory expectations. A single shipment alone can generate up to 50 pages of documentation across 30 parties, and still only about 1% of bills of lading are digital today. Source: McKinsey, “The multi-billion-dollar paper jam: Unlocking trade by digitalizing documentation”. Generic AI can extract information from claims, but extraction is only part of the challenge. The real difficulty is understanding that information in the context of the policy and applying it consistently at scale.

Wording is bespoke, and no two policies read the same

Manuscript clauses, endorsements, exclusions, sub-limits, retroactive dates and layered coverage all have to be read against the specific claim. Nothing about that changes with claim size, the same interpretation work sits behind a cargo shortage and a D&O notification. At high volumes, there is not enough time to review every claim in depth. With complex claims, it is difficult to apply the same level of scrutiny consistently.

Evidence has to be read together, not in isolation

Survey reports, bills of lading, pleadings, forensic reports, expert opinions and financial loss calculations each carry part of the answer. The coverage question is rarely settled by any one of them, and the conflict between two documents is often the finding that matters most.

Claims keep developing after the file is opened

New facts, expert views and litigation positions change coverage, reserves and strategy long after notification. Nothing re-reads the policy against what has arrived since, so the position set at the start is the position the file carries until someone goes back to it.

THE CHALLENGE

From claim data to decision

Claim dimension	What Sprout.ai contextualises	What it returns
Policy bundle	The document payload as it arrives, policy wording, schedules, endorsements and correspondence, across the policy and its renewals	What each document is, the coverages and limits stated, the endorsements in scope, and whether anything needed to reach a judgement is missing
Governing terms	Manuscript clauses, endorsements, retroactive dates and layered coverage, against the date and circumstances claimed	Which wording governs this claim, which endorsement overrides the base terms, and why
Coverage	The governing wording, against the circumstances described and the loss claimed	Whether the claim is covered or excluded, and the clause that says so
Limits and sub-limits	Limits, sub-limits, deductibles and retentions across the layers in place	The limits that apply, and any sub-limit that caps the outcome
Cause and circumstances	Notification data, loss and damage descriptions, and the supporting account of what happened	The cause claimed, what was affected, and where and when it occurred

THE CHALLENGE

From claim data to decision

Claim dimension	What Sprout.ai contextualises	What it returns
Specialist evidence	Survey and forensic reports, expert opinions, pleadings, financial loss calculations and third-party correspondence	The facts that decide the claim, and how the evidence bears on the coverage question
Cross-document consistency	Every document in the file, read against the others	Where accounts agree, where they conflict, and which conflict changes the coverage position
Claim integrity	The parties named on the claim against sanctions and watch lists, and the full file against claim history	Sanctions and screening exposure on any party to the claim, plus duplicates and inconsistencies worth investigating
Completeness	What was submitted, against what the coverage decision needs	Whether there is enough to reach a coverage position, or what is missing and who supplies it
Recommendation	All of the above, reviewed together	The recommended coverage position, the governing clause, and the reasoning that connects them

HOW SPROUT.AI IMPROVES CLAIMS PERFORMANCE

One coverage check, run the same way on every claim in the book

How can specialty insurers give every claim a full coverage check without adding specialists?

How can specialty insurers give every claim a full coverage check without adding specialists?

Sprout.ai runs a set of specialist AI agents on every claim, each assessing a different dimension, and summarises their findings into a single coverage recommendation. The pipeline is built from Sprout.ai's core AI agent building blocks, tuned to the demands of specialty coverage assessment.

The work is the same whether a claim is routine or complex. Teams still need to read the wording, understand which terms apply and assess the evidence against them. What changes is how practical it is to do this manually at scale. Applying the same assessment to every claim removes that trade-off and helps identify routine claims that are becoming more complex while they are still efficient to handle.

Extraction and evidence AI agent

Classifies and reads every document in the file including policy wording, schedules, endorsements, survey and forensic reports, pleadings, correspondence, across formats and languages, and establishes what the claim actually says. Confirms the documents needed to reach a judgement are present, and names what is missing when they are not.

Triage AI agent

Reads the account of the circumstances and establishes what happened, then decides whether there is enough here to reach a coverage position at all. Where there is not, it returns the additional information needed and names the gap rather than making a decision when the evidence is insufficient.

Policy and coverage AI agent

Establishes which wording governs this claim, across renewals, endorsements, retroactive dates and layers, then applies its terms: exclusions, sub-limits, deductibles, and the endorsements that reverse what the base wording would otherwise give. Returns the coverage position and the reasoning behind it.

Evidence analysis AI agent

Reads the specialist evidence against the coverage question rather than document by document: what the survey, the expert opinion and the financial calculation each support, where they conflict, and which of those conflicts changes the position.

Claim integrity AI agent

Screens the parties named on the claim, insureds, counterparties, consignees and intermediaries, against sanctions and watch lists as part of the same pass that reads the file. This means any possible compliance failure surfaces before the claim is worked rather than at payment. Also checks the file against itself and against claim history for duplicates and inconsistencies.

The recommendation

The AI agents' findings resolve into one output per claim: the recommended coverage position, the governing document, the clause applied, the evidence considered, and the reasoning that connects them, traceable back to the document, section and page it came from.

THE KPIS SPROUT.AI MOVES

- Coverage review consistency
- Coverage decision accuracy
- Coverage dispute rate
- Complex claim cycle time
- LAE / ALAE
- Reserve adequacy
- Claims per adjuster
- Decision consistency
- Override rate

65%

of cargo damage comes from poorly packed or secured containers. Sprout.ai checks every claim the same way, so routine ones move fast and complex ones surface early.

Source: TT Club, insurance claims record analysis





30%

of annual marine claims value comes from just the most costly 1% of claims. Sprout.ai's triage AI agent flags high-risk claims before specialist time is committed, not after.

Source: Allianz Commercial / Nordic Association of Marine Insurers

Only 1%

of bills of lading are digital, despite shipments generating up to 50 pages of paperwork. Sprout.ai reads paperwork as it arrives, in any format or language.

Source: McKinsey & Company



FREQUENTLY ASKED QUESTIONS

AI for specialty insurance claims

Q. What happens when a routine claim turns into a complex one?

A. Sprout.ai re-reads the file as evidence arrives and surfaces where the coverage position has moved. That is the point at which a claim being worked as routine needs to be reclassified, and it is normally found late, if at all.

Q. Does AI replace specialist adjusters?

A. No. Sprout.ai gives specialists a cited view of the wording and the evidence, so less of their time goes on assembling the file and more of it on the judgement. The recommendation is theirs to accept, override or escalate, and the record shows which they did.

Q. How does AI handle bespoke policy wording?

A. It reads the wording rather than matching it to a template. Sprout.ai identifies which document and which clause governs the claim such as manuscript clauses, endorsements, retroactive dates and layered coverage, then applies the terms and cites the clause it applied. Nothing has to be encoded in advance, which is why it holds up on wording no two policies share.

Q. How does Sprout.ai support arbitration or litigation defence?

A. Every recommendation records the documents considered, the clause applied, the reasoning, the role of Sprout.ai and the role of the adjuster. That gives you a traceable record of who did what, why and when and is reviewable internally, explainable to an insured or broker, and reconstructable for a regulator or a dispute.

Q. Can it read expert and third-party evidence, not just the policy?

A. Yes. Survey and forensic reports, pleadings, expert opinions and financial loss calculations are read against the coverage question and against each other, so the finding includes where two documents disagree, often the thing that actually decides the claim.

Q. Our documents arrive in several languages and every format imaginable. Does that break it?

A. No. Sprout.ai classifies and reads the payload as it arrives, across formats and languages, rather than requiring a standard submission. That variety is what conventional extraction struggles with, and it is worst at volume.

sprout.ai

About Sprout.ai

Sprout.ai is a London-headquartered global technology company transforming insurance claims with AI-powered automation and decision intelligence. Since 2018, Sprout.ai has partnered with leading insurers, MGAs, and TPAs across major markets to deliver real-time, accurate, fair, and compliant decision-making at scale. Customers include AXA, MetLife, AdvanceCare, and Total Scholastic Solutions (TSS).

With proven and rapid enterprise deployment, Sprout.ai helps insurance teams boost efficiency, strengthen compliance, and deliver exceptional customer experiences. Sprout.ai injects intelligence into insurance claims - changing insurance, for good.

Changing insurance, for good.

Contact our claims experts for more information

Contact us

Book a demo

© 2026 Sprout.ai Limited. All rights reserved.