

GENASYS

THE AI REALITY CHECK

95%

OF UK INSURERS
SURVEYED USE AI

7%

OF INSURERS
HAVE SCALED IT

11%

CAN CLEARLY
SEE A RETURN

Where AI is delivering in underwriting, claims and distribution, the threats it brings and the opportunities it opens for insurers, MGAs and brokers.

genasys.tech.com

Contents

	Introduction	3
I	WHERE AI IS WORKING	
1	Where adoption stands in 2026	5
2	Underwriting and pricing	9
3	Claims and fraud	12
4	Distribution and customer service	16
5	Agentic AI	19
	Case studies	22
II	WHAT IT DEPENDS ON	
6	Data and the core system	24
7	Regulation	26
8	The risks that need managing	30
III	MAKING IT WORK	
9	Building an AI programme that delivers	33
	How Genasys approaches AI	38
	Frequently asked questions	40
	References	42

Introduction

Most insurers already use AI somewhere in the business. Few have scaled it, and fewer still can show what it returns.

This eBook sets out where AI is delivering in insurance in 2026, what it takes to use it well and what that means for insurers, MGAs and brokers in the UK and South Africa, particularly those writing commercial business at £20 million to £250 million of premium.

Its argument is that, for firms of that size, the AI decision with the most effect is the choice of core systems and the state of the data held in them. Three findings run through the evidence. The AI with measured results today mostly reads, extracts and summarises, and outside simple personal lines it leaves decisions that commit money or determine cover with a person. Mid-sized firms will get most of their AI inside the platforms and services they buy, because they will not employ the data science teams that large insurers fund. The rules that govern AI in both markets require a firm to explain and revisit automated decisions, which depends on what its systems record.

Three findings

1

AI reads before it decides

The AI with measured results mostly reads, extracts and summarises, and leaves decisions that commit money or determine cover with a person.

2

It arrives inside platforms

Mid-sized firms will get most of their AI inside the platforms and services they buy.

3

The rules run on records

Explaining and revisiting an automated decision depends on what the firm's systems record.

A firm whose data is captured in structured form, and whose core system exposes that data through APIs, records every decision and lets the business change its own rules, can adopt better AI as it appears. A firm without those foundations will struggle to use AI well whichever tools it buys.

This edition replaces the earlier Genasys whitepaper of the same name. Every figure has been re-sourced to a primary publication or, where the primary was unavailable, to named reporting of it, and the references at the end list each source. Where a figure is a forecast, a consultancy estimate or a technology supplier's account of a client project, the text says so. Since the first edition, generative AI has moved into production at a minority of insurers and attention has turned to agentic systems. The evidence on value has become more sobering. The regulatory position has become clearer, with no AI-specific rules planned in the UK, new rules on automated decisions in force since February 2026 and the EU's high-risk deadline moved to December 2027.

Chapter 1 looks at how far adoption has gone. Chapters 2 to 5 cover underwriting, claims, distribution and agentic AI, followed by six case studies. Chapters 6 to 8 deal with the data and systems AI depends on, the regulation that governs it and the risks that need managing. Chapter 9 sets out a practical approach for a mid-sized firm, with a rating of the main use cases, a readiness checklist and questions to put to suppliers. A short section then describes how the Genasys platform supports AI, and answers to common questions close the book.

Where adoption stands in 2026

The 2025 and 2026 research agrees on one thing above all: using AI, running it at scale and being able to show its return are three different positions, and while most insurers hold the first, only a small minority hold either of the others. The pattern holds in the UK, South Africa and Europe, though the figures behind it come from very different samples.



Bank of England and FCA, 2024 · FSCA and Prudential Authority, 2025 · BCG, 2025 · KPMG, 2026

United Kingdom

The figure most often quoted for the UK, that 95% of insurers use AI, comes from the Bank of England and FCA survey Artificial intelligence in UK financial services, published in November 2024. Insurance reported the highest rate of any sector, against 75% across all respondents. Two caveats apply. Only 19 of the 118 firms that responded were insurers, and they were predominantly large ones. The survey has also not been repeated, so it remains the most recent official measure two years later.

The same survey is more revealing on how AI is used than on how many firms use it. Some 55% of use cases involved a degree of automated decision-making, but only 2% were fully autonomous. A third were third-party implementations, and the three largest providers accounted for 73% of cloud, 44% of model and 33% of data provision. Firms rated 16% of their use cases as high materiality, and those were most common in general insurance. A follow-up Financial Stability in Focus report in April 2025 noted that AI models are widely used in insurers’ pricing and underwriting, and that around half of survey respondents had only a partial understanding of the AI they use.

Share of AI supply held by the three largest providers, UK financial services




Bank of England and FCA, 2024

South Africa

The South African picture is very different. The FSCA and Prudential Authority survey, run between October and December 2024 and published in November 2025, drew around 2,100 responses, of which 220 used AI. As reported by Moonstone, adoption stood at 8% among insurers against 52% among banks and 50% among payment providers, and 41% of insurers planned to spend less than R1 million on AI. Insurers told the regulators they intend to extend AI into underwriting and claims. Across all sectors, the barriers named were POPIA compliance, a shortage of skills and the difficulty of explaining model outputs, while budgets and legacy systems weighed most on smaller firms (Baker McKenzie).

AI adoption by sector, South Africa



FSCA and Prudential Authority, 2025, as reported by Moonstone

The 8% and the UK's 95% measure different populations. The South African base includes thousands of small financial services providers, while the UK figure describes 19 large insurers. Read together, they say that large insurers in both markets use AI and that the long tail of smaller firms mostly does not yet.

From use to scale

International research converges on a small minority of insurers running AI at scale. BCG's September 2025 report Insurance Leads in AI Adoption. Now It's Time to Scale. found that 7% of insurers had scaled AI and about two thirds were still piloting. BCG attributed 70% of the difficulty in scaling to people and process, 20% to technology and 10% to the models themselves. EIOPA's February 2026 survey of 347 European insurers found nearly two thirds using generative AI, most of them still at proof of concept, and 49% with a dedicated AI policy.

Where the difficulty in scaling AI lies



BCG, 2025

The question of return is where the evidence is thinnest. KPMG's Unlocking AI value in insurance, published in September 2026, found that 11% of insurance leaders described their AI return on investment as very clear, and 18% that they had a clear AI strategy. The sample was 44 leaders, so the figures are indicative. McKinsey's cross-industry State of AI 2026, drawn from 1,719 respondents, found 44% of organisations scaling AI but only 37% reporting any effect on EBIT, and the share of high performers unchanged at about 6%. At the top of the market, Evident's AI Index for Insurance 2026 found that 20 of the 30 largest insurers in North America and Europe now report at least one AI use case with measured outcomes, eight more than a year earlier.

One statistic needs handling with care. MIT Nanda's July 2025 report The GenAI Divide claimed that 95% of organisations were getting zero return from generative AI, and the figure has been repeated widely since. It rests on 52 interviews and 153 conference survey responses, and the authors describe their figures as directionally accurate. HSBC's Yuning Bai called the evidence rather weak. This eBook does not use the headline figure.



*Using AI, running it at scale and being able to
show its return are three different positions.*

Where the money goes

Venture investment has moved almost entirely to AI. Gallagher Re's Global InsurTech Report for Q2 2026 found that 99.1% of the quarter's US\$2.44 billion of insurtech funding went to AI-focused companies, after 95% in the first quarter. These figures measure money flowing into technology companies, not what insurers spend.

Insurers' own budgets are a better guide for a mid-sized firm. BCG's research found that insurers spending under US\$5 million a year on AI tend to stay in pilots, while those it classed as leaders spend US\$25 million or more, often US\$50 million to US\$100 million. An MGA or insurer writing £20 million to £250 million of premium will not spend at that level. For firms of that size, most of the AI they use will arrive inside platforms and services they already buy, which makes the choice of those platforms an AI decision as much as an operational one.

What customers think

UK consumers are open to AI that helps and wary of AI that decides. Research by Yonder for the FCA, carried out in April 2026 with 5,026 adults and published alongside the Mills Review, found that 26% regard AI as a reliable source of financial information. Willingness to use AI fell from 36% for assistive tools to 30% for AI acting with the customer's permission and 20% for AI acting autonomously. Only 40% knew they would have no formal recourse if they acted on general AI advice. In South Africa, the KPMG and University of Melbourne trust study found 62% of people willing to trust AI and 67% believing it should be regulated.

UK adults willing to use AI in financial services, by how much it does



FCA and Yonder, 2026. 5,026 adults.

The surveys behind the headline figures

SOURCE	PUBLISHED	SAMPLE	HEADLINE FINDING
KPMG, Unlocking AI value in insurance	Sep 2026	44 insurance leaders	11% say AI return is very clear
McKinsey, State of AI 2026	Aug 2026	1,719 organisations, all sectors	44% scaling AI; 37% report EBIT impact
Evident, AI Index for Insurance	Jun 2026	30 largest insurers, North America and Europe	20 of 30 report measured outcomes
EIOPA, generative AI survey	Feb 2026	347 European insurers	Nearly two thirds use GenAI; most at proof of concept
FSCA and Prudential Authority	Nov 2025	About 2,100 SA financial firms	8% of insurers use AI
BCG, Insurance Leads in AI Adoption	Sep 2025	Insurers (sample not stated)	7% have scaled AI
Bank of England and FCA	Nov 2024	118 UK firms, 19 insurers	95% of insurers use AI


Underwriting and pricing

The measured gains in underwriting so far have come from the work around the decision: reading submissions, triaging them and finding the right guidance. Fewer insurers have published results for AI making or materially changing the underwriting decision itself, and those that have tend to describe it in qualitative terms.

Where underwriters' time goes

The case for AI in underwriting starts with how little of an underwriter's day is spent underwriting. A 2021 Accenture survey, reported in 2022, found underwriters spent 70% of their working time on other activity, with 40% on administration and 30% on negotiation and sales support. No later survey has replaced that figure. Accenture's August 2025 research, *Underwriting Rewritten*, surveyed 430 senior underwriting executives, 140 of them in commercial property and casualty, and found they expected AI's effect on underwriting tasks to rise from 17% today to 75% within three years.

How an underwriter's working time is spent



Accenture, 2022. Underwriting share is the remainder.

In commercial lines the administrative load sits mostly in submission intake. A broker submission typically arrives as an email with a proposal form, a schedule of values, loss runs and other attachments in inconsistent formats, and someone has to read, rekey and check it before an underwriter can assess the risk. This is the step where AI has the clearest record. Zurich has rolled out an AI intake and triage platform across its commercial underwriting operations, going live in five countries within 90 days. Figures released with the platform's provider, Cytora, in May 2026 put the reduction in manual triage time at 80%, from 75 minutes to 15, and the share of submissions processed straight through at intake at 95%, up from 10%. The figures come from a vendor announcement and have not been independently audited.

Zurich commercial intake, before and after

MANUAL TRIAGE TIME PER SUBMISSION



SUBMISSIONS PROCESSED STRAIGHT THROUGH



Zurich and Cytora, May 2026 (supplier account)


Guidance retrieval is a smaller and more easily replicated gain. AXA built a retrieval tool that lets underwriters query internal guidelines in plain language, and reported that it cut the time to find the relevant guidance from ten minutes to under three, with 86% of pilot users rating it eight out of ten or higher.

AI in the underwriting decision

Hiscox's work with Google at Lloyd's is the most cited example of AI contributing to a lead underwriting decision. When Hiscox announced the project in December 2023, it said AI could cut the time for a lead open-market quote from three days to three minutes. By August 2024 the model was live for sabotage and terrorism renewals in the US and Canada, producing quotes for brokers within minutes. Hiscox's 2025 results describe it as having pioneered AI-augmented underwriting at Lloyd's, without publishing performance figures. The application is deliberately narrow: a single class of business with standardised submissions and a human underwriter retaining authority.

Pricing is the area where AI has been in production longest, particularly in personal lines. The Bank of England's April 2025 review noted that AI-based models are widely used to support insurers' pricing and underwriting decisions. Aviva's 2025 results refer to pricing sophistication delivered through AI models, again without figures. Commercial lines have less to work with: fewer policies, more heterogeneous risks and thinner claims histories mean that the volume of clean data a pricing model needs is harder to assemble.

Underwriters' confidence is a limiting factor. Capgemini's 2024 World Property and Casualty Insurance Report found that 43% of underwriters trusted and regularly accepted the recommendations of predictive tools, 59% were concerned about the integrity of the underlying data and 73% said weak data limited pricing accuracy.



Underwriters surveyed by Capgemini, 2024

Commercial lines, Lloyd's and MGAs

The Lloyd's market has moved quickly from a low base. A survey by the LMA and Barnett Waddingham, published in April 2026, drew 39 responses covering more than 60% of Lloyd's stamp capacity. About a quarter of respondents used AI in 2025 and adoption had more than doubled a year later. Some 72% had an AI governance framework in place and a further 21% were building one, and more than 60% required human review of AI outputs. Use in core underwriting, pricing and claims remained small.

The commercial insurers that have published results report gains in the same places as Zurich and AXA. Allianz UK gave all its commercial property and liability underwriters a generative AI assistant, known as BRIAN, that answers questions on underwriting guidance; by September 2025 they had asked it more than 13,000 questions and saved around 65,000 minutes, according to Claims Magazine. Markel reported in 2023 that AI-assisted submission handling, also supplied by Cytora, had cut quote turnaround from a day to two hours.

MGAs expect AI to support decisions more than make them. In research by Clyde & Co with the MGAA, published in July 2025, 65% of MGAs and 66% of their capacity providers said they expected AI to play a supporting role in decision-making, and 77% of MGAs said their claims handling needed to improve, up from 59% in 2023. With more than 350 MGAs in the UK writing over a tenth of the general insurance market (Insurance Business, July 2026), the way MGAs adopt AI will shape a large share of commercial lines.


South African underwriting managers describe the same administrative load. Speaking at the SAUMA conference in June 2025, Deen Pillay of Emerald UMA put the share of an underwriter's time spent on administration at 80% (FAnews). Santam's 2025 results presentation lists using data and AI to improve pricing and risk selection as a priority, without published results so far. At the automated end of the market, Naked sells all of its car, home and single-item policies online without human intervention and produces quotes in under 90 seconds (Creamer Media, January 2025).

The size of the prize

Consultancy estimates of underwriting value are modelled rather than measured, and should be read that way. BCG's March 2026 analysis of AI-first property and casualty insurers, built on US market data, estimates a 30% to 40% reduction in underwriters' active handling time, a similar reduction in time to quote and a 15% to 20% reduction in the underwriting expense ratio.

For an MGA or insurer writing commercial business, the practical sequence follows from the evidence. Submission intake and triage have the most proven results, the lowest regulatory exposure and the fastest payback, because the underwriter still makes the decision. AI that informs the decision comes next and depends on the firm holding structured exposure and claims data it can trust. AI that makes the decision is, outside narrow and standardised classes, still rare.

The practical sequence



Claims and fraud

Claims is where insurers have published the most concrete AI results, and fraud is where AI now works for both sides. UK insurers detected £1.34 billion of fraudulent claims in 2025, the highest value on record, while reporting a growing number of claims supported by AI-generated images and documents.

Claims handling

Aviva has gone further than most UK insurers in applying AI to claims. A McKinsey case study of the programme, which McKinsey supported, reports more than 80 AI models in the claims journey, a 23-day reduction in the average time to assess liability on complex cases, a 30% improvement in routing accuracy and a 65% fall in customer complaints. Aviva's own 2025 results refer to significant claims indemnity benefits from AI without giving a figure, so the McKinsey numbers are the most specific available and come from the adviser rather than the insurer.



Aviva, as reported by McKinsey (adviser's account)

Allianz has published the clearest account of an agentic claims process. Project Nemo, which went live in Australia in July 2025, handles food spoilage claims under AUD 500 using seven AI agents that cover planning, cyber checks, coverage, weather verification, fraud, payout calculation and audit. Allianz reports an 80% reduction in processing and settlement time, with claims reaching a human reviewer within five minutes of filing. The system was built in under 100 days, and Allianz states that payout decisions are never automated by design.

Project Nemo: seven agents, one human decision



Allianz, November 2025. Food spoilage claims under AUD 500.

At the far end of automation, Lemonade’s 2025 annual report states that its AI takes the first notice of loss in 96% of claims and that roughly 55% of claims were fully automated. Lemonade writes US personal lines on a platform built for automation from the start, so its figures show what is possible in simple, high-volume products and say little about commercial claims with multiple parties, coverage questions and loss adjusters.

Consultancies have estimated the wider value. Bain’s October 2024 analysis projected a 20% to 25% reduction in property and casualty loss-adjusting expenses and a 30% to 50% reduction in claims leakage from generative AI. BCG’s March 2026 estimate put the reduction in loss-adjusting expenses at 20% to 30%. Both are modelled projections.

Fraud in the UK


The ABI’s figures for 2025, published in October 2026, show insurers detected 93,900 fraudulent claims worth £1.34 billion. The value was 14% higher than in 2024 (£1.18 billion, revised from the £1.16 billion first reported), although the number of cases fell by 2.7%. The average fraudulent claim was £14,300. Motor accounted for 51,900 cases worth £625 million, and insurers stopped a further £1.8 billion of fraud at application. The ABI’s 2024 release had already described fraudsters taking more sophisticated approaches aided by AI.



Association of British Insurers, October 2026

The value of detected fraud fell from £1.3 billion in 2015 to £1.1 billion between 2020 and 2022 and has risen since, while the number of cases in 2025 was around 28% lower than in 2015. The average detected fraudulent claim has grown accordingly.

Value of detected fraudulent claims, UK (£bn)



ABI annual fraud releases and reporting of them. Selected years; values before 2023 rounded to £0.1bn.

The largest insurers’ own figures point the same way. Aviva reported motor fraud by value up 39% and a growing number of claims supported by AI-generated images and manipulated documents. Allianz UK reported staged accidents up 84% and said it uses voice analytics and machine learning to detect them. Allianz had earlier reported a 300% rise between 2022 and 2023 in cases where apps had been used to alter images, videos and documents. At the point of sale, the Insurance Fraud Bureau reported that ghost broking had risen by more than half in two years.

INSURER	PERIOD	SUSPECT OR DETECTED CLAIMS	VALUE	SOURCE
Aviva (including Direct Line)	2025	More than 18,400	£233m	Aviva, June 2026
Allianz UK	2025	More than 34,200	Almost £174m	Allianz UK, March 2026
Zurich UK	2024	4,049	£94.8m	Zurich UK, March 2025

The industry’s readiness for AI-generated fraud is low. A 2026 survey by the Association of Certified Fraud Examiners and the analytics firm SAS found that 7% of anti-fraud professionals considered their organisation more than moderately prepared for it.

*Any process that accepts images or documents as evidence
now needs a way to test whether they are genuine.*

Fraud in South Africa

South Africa has no current, industry-wide measure of short-term insurance fraud. The R4 billion a year still quoted in the press dates from 2013. The Insurance Crime Bureau has estimated that the industry could lose R3.5 billion if fraud and organised crime go unchecked, based on 10% of the R35 billion paid in short-term claims each year, so the figure describes potential exposure. In life insurance, ASISA members recorded 16,520 cases of fraud and dishonesty in 2024 and prevented around R1.4 billion of losses. Santam has described using AI for image recognition and deepfake detection in claims, without publishing results.

Several South African insurers have described AI in their claims operations, mostly without figures. King Price told ITWeb in July 2026 that an AI claims risk engine approves routine claims automatically and has halved their turnaround, and that it uses voice identification to detect misrepresentation and is investigating deepfake image detection. In the same article Discovery Life said AI and analytics flag suspicious claims and that fraud accounts for less than 0.5% of its claims. Momentum Insure said in November 2025 that it was putting AI at the core of its claims management, and Old Mutual Insure has warned customers about AI-generated images submitted as claims evidence.

The Insurance Crime Bureau says it uses AI alongside trained investigators and has prevented or recovered R5.4 billion over a decade (FAnews, June 2026). Its chief executive, Garth de Klerk, has warned that criminals are using AI to increase the volume of attacks on insurers and that people are being caught out by deepfake videos. No South African body yet publishes figures on AI-enabled claims fraud.

What this means for a commercial lines book

For a mid-sized commercial insurer or MGA, the transferable lessons from these programmes are about design more than scale. Allianz automated the assembly of evidence and kept the payment decision with a person. Aviva applied many narrow models to specific steps rather than one model to the whole claim. Both depended on claims data being captured in structured form at first notification, which is where many commercial claims processes still rely on email, PDFs and a loss adjuster's report. The fraud figures add a further requirement: any process that accepts images or documents as evidence now needs a way to test whether they are genuine.

Distribution and customer service

Customers accept AI that gets them to an answer faster and resent AI that stands between them and a person. The FCA's Consumer Duty work takes the same position, and Klarna, the best-known company to replace service staff with AI, has started to reverse.

What customers will accept

A Gartner survey of 5,728 customers, published in July 2024, found that 64% would prefer companies did not use AI in customer service and 53% would consider switching to a competitor that did not. Their main concern was that AI would make it harder to reach a person. The FCA research described in Chapter 1 found the same pattern in financial services, with willingness to use AI falling as the AI is given more autonomy.

**64%**

PREFER COMPANIES DID NOT USE AI IN SERVICE

**53%**

WOULD CONSIDER SWITCHING TO A COMPETITOR THAT DID NOT

Gartner survey of 5,728 customers, July 2024

Klarna is the best-documented example of what happens when a firm moves too far. In February 2024 the company said its AI assistant was handling two thirds of customer service chats, the work of 700 full-time staff. In May 2025 its chief executive told Bloomberg that the company was recruiting human agents again and that investing in the quality of human support was its future direction. Gartner predicted in June 2025 that half of the organisations planning significant cuts to customer service staff because of AI would abandon those plans by 2027.

The FCA does not prescribe channels under the Consumer Duty, but it expects the channels a firm offers to meet customers' needs. Its review of the consumer support outcome cited as good practice firms that added human contact wherever complex needs were identified, including routing bereavement queries from a chatbot to a person. For an insurer, a claim is often the moment of greatest need, so an AI-led claims channel has to make the route to a human obvious and quick.

*A claim is often the moment of greatest need, so
the route to a human has to be obvious and quick.*

Where AI is working in service

The service deployments with the clearest results keep a route to a person open. Travelers, the US insurer, uses an autonomous voice assistant built with OpenAI to take motor claims. According to OpenAI's account of the deployment, 85% to 90% of callers who start with the assistant complete their claim with it, and callers can ask for a live claims specialist. Lemonade, as Chapter 3 noted, handles more than half of customer enquiries through AI on a platform designed for it.

In commercial lines the customer that service teams deal with most is the broker. Brokers want quotes, mid-term adjustments, documents and claims updates, most of which are retrieval and assembly work that AI handles well when the underlying policy data is structured and accessible.


Brokers

The largest broking groups are now spending at a scale that will change how they deal with insurers. Marsh's internal assistant, LenAI, is available to staff in more than 80 countries and was handling about a million prompts a week by August 2026, and the firm has built around 130 digital agents (Insurance Business, August 2026). WTW announced in July 2026 a programme costing around US\$625 million to embed AI and automation across its operations by the end of 2028, targeting about US\$400 million of annual run-rate savings (Insurance Business, July 2026). Gallagher has rolled out an AI platform to around 60,000 staff and, in a Microsoft case study, reports claims review time cut by up to 90%. Acrisure has announced plans to cut around 2,250 jobs, about 11% of its workforce, in a restructuring tied to AI (Business Insurance, June 2026).



Smaller brokers are a long way behind. In a survey of 207 broking firms by Open GI, a broker software provider, reported in May 2025, 80% had not implemented AI. Some 45% of national brokers had done so against 9% of regional and provincial firms, and 54% of smaller brokers named integration with their existing systems as a barrier.

UK brokers that have implemented AI



Open GI survey of 207 broking firms, as reported by Insurance Age, May 2025

South African insurers have taken opposite positions on automating service. OUTsurance's chief executive, Marthinus Visser, said in September 2025 that the personal touch still matters, and the company recruited nearly 1,000 additional call-centre staff that year. Naked has gone the other way and launched a ChatGPT app that gives final car insurance quotes. Sanlam appointed a group chief AI officer in March 2026 (Sanlam).

For insurers and MGAs, broker AI changes what good service looks like. A broker whose own systems can prepare a submission in minutes will expect an answer from the market at a similar speed, and will route business towards insurers whose systems can receive structured data, return quotes and documents by API and answer servicing questions without a phone call.


Agentic AI

Agentic AI describes systems that take a sequence of actions towards a goal, using tools and data, instead of answering a single prompt. A small number of insurers have agentic systems in production on narrow, bounded processes. Gartner estimates that most products marketed as agentic fall short of that definition and expects a large share of current projects to be abandoned.

What the term means

The Bank of England's April 2025 review describes agentic AI as systems that act autonomously towards set goals, use tools, learn from feedback and adapt to their environment. In practice an insurance agent might read a claim, check the policy wording, pull weather data, query a fraud model, calculate a settlement and draft a decision for a person to approve. The difference from earlier automation is that the system decides the order of steps and which tools to use, which makes it more flexible and harder to test.

An agent at work on a claim



Where it is in production

The published examples are few and deliberately constrained. Allianz's Project Nemo, described in Chapter 3, uses seven agents on food spoilage claims under AUD 500 and leaves every payout decision with a person. Zurich's commercial intake platform, described in Chapter 2, uses agents to read and triage submissions before an underwriter sees them. AXA has said it has more than 60 agentic use cases in testing or partial deployment across contact centres, underwriting and claims, built on a group-wide platform.

Across the largest insurers, agentic systems are still a small share of what is deployed. Evident's AI Index for Insurance 2026 found that 8% of the use cases disclosed by the 30 largest insurers in North America and Europe pointed towards agentic reasoning. KPMG's September 2026 survey found 29% of insurance leaders running some processes end to end through agents or automation, though from a sample of 44.

The case for caution

Gartner predicted in June 2025 that more than 40% of agentic AI projects would be cancelled by the end of 2027 because of rising costs, unclear value or weak risk controls. It estimated that only about 130 of the thousands of vendors marketing agentic products offered genuine agentic capability, calling the rest agent washing, and observed that many use cases presented as agentic did not need an agentic design. The same research forecast that by 2028 at least 15% of day-to-day work decisions would be made autonomously, up from none in 2024.




UK regulators are watching closely. In a January 2026 speech, the FCA's Sheldon Mills said autonomous AI was coming into view and asked what the reasonable steps expected of senior managers under SM&CR mean when a model can change weekly. The Bank of England's April 2026 letter to the Treasury Committee said the Financial Policy Committee had asked for further work on agentic AI and that the risks were likely to increase, potentially rapidly. The FCA's July 2026 Mills Review, summarised in Chapter 7, proposed a principles-based framework for AI agents and kept firms accountable for their outcomes.

Where agents fit in an insurance operation

The deployments that work share four characteristics. The process is high in volume and low in value per item, so an error is cheap to correct. The steps are well defined and each tool the agent calls has a clear, testable output. Every action is logged so that the firm can reconstruct what the agent did and why. A person approves anything that commits money or determines cover. Food spoilage claims, submission intake, document chasing, bordereaux validation and renewal preparation fit that profile. Coverage decisions on a complex commercial claim, a declinature or a large-loss settlement do not, and putting an agent in charge of them would leave the firm unable to explain a decision to the customer, the FCA or a court.

Four characteristics of agent deployments that work



Fits the profile

Food spoilage claims

Submission intake

Document chasing

Bordereaux validation

Renewal preparation

Does not

Coverage decisions on a complex commercial claim

A declinature

A large-loss settlement

Agents also depend on the systems around them. An agent can act only through the interfaces it is given, so an insurer whose policy, claims and billing data sit behind screens with no API has nothing for an agent to work with beyond what it can scrape, which is neither reliable nor auditable.

Case studies

The six case studies below were chosen because they involve specialist writers, MGAs, claims administrators or mid-sized insurers and report measured results. For each, the source of the figures is stated: some come from the firm itself, others from the technology supplier with the firm quoted. No UK MGA or South African commercial insurer had published measured before-and-after results for AI by October 2026, which is why the closest MGA example is from the US.



1 Davies Group: claims intake

Davies, the UK claims and insurance services group, handles claims on behalf of insurers and corporate clients. In April 2026 it introduced ClaimPilot, an AI agent that takes in new corporate risk claims, reads the accompanying documents and sets up the claim for a handler. Davies reported that the agent cut intake process times by more than 80%. The figure is Davies's own; the baseline time was not published.

2 Paragon Insurance Holdings: submission triage at an MGA

Paragon is a US specialty MGA running around 25 programmes. Its highest-volume programmes each received more than 50,000 submissions a year, and its underwriters had capacity to review about 30% of what arrived, so most submissions were never reviewed. Paragon deployed an AI underwriting platform from Kalepa to read, extract and prioritise every submission. According to the case study, in which Paragon's president of excess and surplus lines and its chief technology officer are quoted, the whole inbox was being processed within 90 days at 98% to 99% extraction accuracy, and the quote-to-bind ratio doubled within a year. The account is published by the supplier.

3 Arch Insurance: intake to decision

Arch writes specialty lines in the US and UK, including executive assurance, professional liability and excess casualty. Using an intake platform from Cytora, Arch reduced the time from a submission arriving to an underwriting decision by 70%, from two to three days to around two hours, and moved from quoting in three to four days to quoting the same day, with 95% accuracy in digitising submissions. The figures appear in an article by Cytora's chief executive, May 2026, which quotes Arch on the effect on conversion.

4 Markerstudy: claims calls

Markerstudy, a UK insurance group focused on motor, used Microsoft's Azure OpenAI service to summarise claims calls automatically in place of handlers writing notes. It reported a saving of about four minutes per call and, across roughly 840,000 claims calls a year, an estimated 56,000 hours. The group chief executive is quoted in Microsoft's account; the annual figure is an estimate.

5 Esure: contact centre handling time

Esure, the UK personal lines insurer, applied generative AI to support contact centre staff during and after calls. Its chief data officer told Insurance Post in September 2024 that the tools had cut average handling time by up to 3.5 minutes per interaction, and that the approach could save up to 8% of frontline operating costs. The cost figure is a projection.

6 Hollard Insurance Australia: claims file review

Hollard's Australian business piloted AI to help claims consultants review claim files before making decisions. It reported in March 2026 that review time fell by 70% in the pilot and that consultants saved up to 25 to 35 minutes on complex files. Hollard Australia is a separate business from Hollard in South Africa, and the results are from a pilot.

What the case studies have in common

All six apply AI to reading and summarising: submissions, claim documents, call recordings and claim files. In each, a person still made the underwriting or claims decision, and the measured gain was in the time it took that person to get to it. Paragon is the only case that reports a commercial outcome, a higher quote-to-bind ratio, and it came from being able to look at every submission instead of a third of them.



In each, a person still made the decision, and the measured gain was in the time it took that person to get to it.

Data and the core system

The limiting factor for most insurers’ AI programmes is their data and the systems that hold it. The surveys in this eBook that ask insurers what stops them scaling AI consistently put data quality and legacy systems at or near the top.

What the surveys say

In the Bank of England and FCA 2024 survey, four of the five risks firms ranked highest related to data, led by data privacy and protection, data quality and data security. Capgemini’s World Property and Casualty Insurance Report 2026 found that 12% of insurers described their data maturity as very high, 10% were scaling AI successfully and 60% were stuck at exploration or proof of concept. It also found that 72% of AI spending went on technology and 28% on the organisational change needed to use it. KPMG’s September 2026 survey found 11% of insurers with the data foundations and governance to scale AI beyond pilots. In EY’s spring 2025 survey of 100 US insurers, 68% named data privacy, quality or security as a barrier and 66% named integration with legacy systems.

Where AI money goes, and where the difficulty lies



Capgemini, 2026 · BCG, 2025. Chapter 9 returns to this gap.

Gartner has put a number on the consequence. It predicted in February 2025 that through 2026 organisations would abandon 60% of AI projects that were not supported by AI-ready data, and found that 63% of organisations either lacked the right data management practices for AI or did not know whether they had them.

Market data standards

The London market’s attempt to fix its data at market level has changed shape. Lloyd’s announced on 19 March 2026 that its Blueprint Two programme would be sunset, while saying it remained committed to moving the market onto resilient, cloud-based infrastructure. The data standard developed alongside it has continued. The Core Data Record now covers the full insurance lifecycle, including treaty, claims notification and delegated authority contracts such as binders and lineslips, aligned to ACORD standards. The LMA’s first-quarter 2026 update recorded that only 43% of premium payments were made on time, a measure of how much manual reconciliation the market still carries.

Delegated authority is where the data problem is most acute for MGAs and their capacity providers. Bordereaux arrive monthly in spreadsheets of varying structure, and the market's own data service, LIMOSS, notes that their quality and completeness can vary greatly. An AI model trained on, or making decisions from, inconsistent bordereaux will reproduce the inconsistency at speed.



*An AI model trained on inconsistent bordereaux
will reproduce the inconsistency at speed.*

What AI needs from the core system

The policy administration system is where most of an insurer's structured data is created: the risk details, the rating factors, the cover, the premium, the endorsements and the claims history. Whether AI can be used safely and at scale depends to a large extent on five properties of that system.

- 1 Structured data at the point of entry**, so that a model receives defined fields with validation, and an extraction step is not needed to recover what a person typed into free text or attached as a PDF.
- 2 Open APIs and event notifications**, so that a model can be called at the right step in a process and its output written back to the policy or claim record without a person rekeying it.
- 3 A complete audit trail** that records which model or rule produced each output, what data it used and who approved the result. Consumer Duty, Articles 22A to 22D of the UK GDPR and POPIA section 71 all require a firm to explain and, where challenged, revisit decisions.
- 4 Business rules and workflows that the business can change through configuration**, so that a human approval step, a referral threshold or a new data field can be added without a development project.
- 5 Clean separation between the system of record and the AI services that use it**, so that a model can be replaced as better ones appear without changing the core platform.

Insurers on older platforms can meet some of these requirements with integration layers and data warehouses, at a cost in time and maintenance. For firms considering a change of core system in the next few years, these five properties are a practical test of whether a platform will support the AI the business will want to use.

Regulation

No regulator in the UK or South Africa has written AI-specific rules for insurers. The FCA has said it does not plan to, and the FSCA is developing a framework but had issued no guidance by October 2026. What applies instead is the existing conduct, data protection and accountability regime, which already reaches most of the ways an insurer might use AI. The one hard deadline that UK firms had been planning around, the EU AI Act’s high-risk regime, moved in July 2026 from August 2026 to December 2027.

The regulatory calendar

●	Oct 2024	FCA	AI Lab opens
●	Nov 2025	FSCA AND PA	Joint report on AI in the South African financial sector
●	Dec 2025	FCA	First AI Live Testing cohort, including Homeprotect and Scottish Widows
●	5 Feb 2026	UK	UK GDPR Articles 22A to 22D on automated decisions in force
●	Feb 2026	FCA	Insurance Regulatory Priorities: AI in underwriting, claims and service under review
●	Apr 2026	SOUTH AFRICA	Draft National AI Policy withdrawn
●	Jul 2026	FCA	Mills Review on AI and retail financial services
●	27 Jul 2026	EU	Digital Omnibus on AI in force
⋮	1 Jan 2027	COLORADO	SB 26-189 takes effect
⋮	Jan 2027	SOUTH AFRICA	Revised National AI Policy expected for comment
⋮	Early 2027	FCA	Delegated authority oversight review findings
⋮	2 Dec 2027	EU	AI Act high-risk obligations apply

Sources as cited in this chapter; dashed markers are future dates, and early 2027 dates are approximate

United Kingdom

The FCA’s published approach to AI, last updated in October 2026, says it does not plan to introduce extra regulations for AI. It relies on the Consumer Duty, which requires firms to deliver and evidence good customer outcomes, and on the Senior Managers and Certification Regime, which places accountability for a system’s behaviour with a named individual. An insurer cannot attribute a poor claims outcome to its model. The outcome belongs to the firm and the accountability belongs to a senior manager.

An insurer cannot attribute a poor claims outcome to its model.

The regulator has been specific about where it is looking. Its Insurance Regulatory Priorities report, published in February 2026, says the FCA supports firms' use of AI but expects them to monitor consumer outcomes closely, and that it will assess how firms use AI in underwriting, claims and customer service. The same report confirms investigations into home and travel claims handling following the Which? super-complaint, continuing fair value work and a review of outsourced claims and delegated authority oversight, with findings due in early 2027. For capacity providers working through MGAs, that last review is likely to bring AI used anywhere in a delegated claims or underwriting chain within the oversight the FCA expects them to evidence.

The FCA has also built routes for firms to test AI under supervision. The AI Lab opened in October 2024. The Supercharged Sandbox, run with Nvidia, took its first proofs of concept from October 2025. The first AI Live Testing cohort, announced in December 2025, included the home insurer Homeprotect and the life and pensions provider Scottish Widows. The second cohort, announced in April 2026, had no insurance use case, and the FCA expects to publish its evaluation in the first quarter of 2027. In July 2026 the Mills Review on AI and retail financial services described services moving towards AI-enabled, continuous and delegated models and made seven recommendations to the FCA Board, none of them new rules.

The change with the most immediate effect on day-to-day processing came from data protection law. Section 80 of the Data (Use and Access) Act 2025 replaced Article 22 of the UK GDPR with Articles 22A to 22D and came into force on 5 February 2026. Significant decisions made solely by automated means are now permitted on a wider range of lawful bases, provided the firm tells the individual, lets them make representations, offers human intervention and allows the decision to be contested. Decisions based on special category data, which includes health information used in underwriting and claims, remain restricted. The ICO consulted on draft guidance between 31 March and 29 May 2026; final guidance had not been published at the time of writing.

On the prudential side, the PRA's SS1/23 model risk management principles cover AI in modelling but apply to banks and building societies with internal model approval, not to insurers. In an April 2026 letter to the Treasury Committee, Deputy Governor Sarah Breeden confirmed the Bank has no plans for AI-specific rules and that AI is a PRA supervisory priority for 2026. The government's position is unchanged from its 2023 white paper: five principles (safety, transparency, fairness, accountability and contestability) applied by existing regulators, and no AI bill had been introduced as of September 2026.

European Union

The EU AI Act classes as high-risk any AI system used for risk assessment and pricing of natural persons in life and health insurance (Annex III, point 5(c)). Commercial property and casualty lines are outside that category. The high-risk obligations were due to apply from 2 August 2026, but the Digital Omnibus on AI, adopted as Regulation (EU) 2026/1744 and in force since 27 July 2026, moved the date for Annex III systems to 2 December 2027. The Act applies to providers and deployers outside the EU where the system's output is used in the EU, so a UK group writing life or health business through an EU entity remains in scope.

For AI that falls outside the high-risk category, EIOPA's opinion on AI governance, published in August 2025, sets out how existing Solvency II and Insurance Distribution Directive duties apply to data governance, record keeping, fairness, explainability and human oversight. EIOPA states that the opinion creates no new requirements.

South Africa

The FSCA and the Prudential Authority published their joint report, Artificial Intelligence in the South African Financial Sector, on 24 November 2025. It reports AI adoption of 52% among banks and 50% among payment providers, says sector guidance on ethical AI is envisaged and states that institutions should disclose when AI is used in decisions that affect consumers. No FSCA guidance on AI had been issued by October 2026; the FSCA's integrated report, tabled in September 2026, says an AI framework is being developed.

The binding constraint in the meantime is section 71 of the Protection of Personal Information Act. A person may not be subject to a decision with legal consequences based solely on automated processing intended to profile them, unless the decision is made in connection with a contract or authorised by law or a code of conduct. Where the contract exception is used, the insurer must allow the customer to make representations and must give them enough information about the logic of the automated processing.

Two pieces of wider reform remain unfinished. The Conduct of Financial Institutions Bill, which will embed Treating Customers Fairly in statute, had its notice of introduction published on 17 April 2026 but had not been formally tabled in Parliament by August 2026, and insurance falls in the third phase of its rollout. The draft National AI Policy, approved by Cabinet in March 2026, was withdrawn on 26 April 2026 after Communications Minister Solly Malatsi found it contained fictitious citations. A revised draft is expected for public comment in January 2027.

United States

UK and South African insurers with US exposure should know two developments. The NAIC's model bulletin on insurers' use of AI systems, issued in December 2023, had been adopted by 29 jurisdictions by October 2026. Colorado, which has gone furthest, amended its governance regulation 10-1-1 from October 2025 to cover life, private motor and health insurers, and replaced its broader AI Act with SB 26-189, which takes effect on 1 January 2027 and lists insurance among the decisions it covers.

JURISDICTION	AI-SPECIFIC RULES FOR INSURERS	WHAT APPLIES NOW	DATES TO NOTE
UK	None planned	Consumer Duty, SM&CR, UK GDPR Articles 22A to 22D	ICO automated decision-making guidance pending; delegated authority review findings early 2027
EU	Yes, for life and health pricing of individuals	AI Act; EIOPA opinion under Solvency II and IDD	High-risk obligations from 2 December 2027
South Africa	None yet	POPIA section 71; TCF	Revised National AI Policy for comment January 2027; COFI Bill awaiting tabling
US	Varies by state	NAIC model bulletin in 29 jurisdictions	Colorado SB 26-189 from 1 January 2027

What the rules require of systems

The rules differ by market, but what they ask of an insurer's systems is consistent. Under the Consumer Duty a firm has to evidence the outcomes its customers receive. Under SM&CR a named senior manager has to be able to show the reasonable steps taken to control a model. Under Articles 22A to 22D of the UK GDPR and section 71 of POPIA, a customer affected by an automated decision can ask for the logic to be explained and for a person to review it, and under the EU AI Act a high-risk system must keep logs. Each of those duties is met from records: which model or rule produced an output, what data it used, who approved it and what happened next. A firm whose core systems keep that record can meet them for every AI service it adds, while a firm that has to reconstruct it from emails and spreadsheets will find each new use of AI harder to defend.

The risks that need managing

The risks of AI in insurance are specific and documented: unfair outcomes for groups of customers, decisions the firm cannot explain, outputs that are confidently wrong, dependence on a few suppliers, fraud that uses the same tools and disruption to the workforce. Each has a recent example, and each falls on the firm and its senior managers rather than on the technology provider.

Bias and proxy discrimination

Insurers price on risk, and some risk factors correlate with characteristics the law protects. In March 2022 Citizens Advice published mystery shopping research across 649 quotes showing that people in postcodes with a high proportion of people of colour were quoted at least £280 a year more for car insurance, and that crime, deprivation and accident rates did not explain the difference. Its 2023 follow-up put the gap at around £250 a year and the total cost at £180 million in 2022. Insurers responded that they do not collect ethnicity data, which is the point: a model can produce a discriminatory outcome without ever seeing the protected characteristic, by learning from variables that stand in for it.

Correcting bias is harder than detecting it. The FCA's research note on bias in supervised machine learning, first published in December 2024, identifies past decisions and historical exclusion as the main sources of bias in training data, and warns that mitigation can reduce a model's accuracy or worsen outcomes for a different group. A firm therefore needs to decide in advance which measure of fairness it is testing against and to record why.

Explainability and accountability

As Chapter 1 noted, around half the firms in the Bank of England's 2024 survey said they had only a partial understanding of the AI they used. That position is difficult to sustain under the Consumer Duty, which requires a firm to evidence good outcomes, and under the UK and South African rules on automated decisions described in Chapter 7, which require it to explain the logic of a decision to the person affected.

Tribunals have already held companies to their AI's statements. In February 2024 the British Columbia Civil Resolution Tribunal ordered Air Canada to honour a refund policy its website chatbot had described incorrectly, rejecting the airline's argument that the chatbot was responsible for its own information. The sum was about CAD 650. The principle applies to any insurer whose chatbot explains cover or claims procedures.

Outputs that are wrong

Generative models produce fluent text whether or not it is accurate. In October 2025 Deloitte Australia agreed to partially refund the Australian government for an AUD 440,000 report that contained fabricated references and a fabricated quotation attributed to a Federal Court judgment; Deloitte had disclosed using a generative AI model in its preparation. The South African government's withdrawal of its draft National AI Policy in April 2026, described in Chapter 7, followed the discovery of fictitious citations. In insurance the equivalent failures would be a policy summary that misstates an exclusion, a claims letter that cites the wrong clause or a broker response that quotes a premium the rating engine never produced. Each of them is a conduct issue, and each is prevented by grounding the model in the firm's own records and checking its output before it reaches the customer.

Concentration and third parties

Chapter 1 showed how much of the industry's AI comes from a small number of third-party providers. The Bank of England's April 2025 review warned that growing concentration in the supply of AI services could increase risk to the financial system. For an individual insurer, the practical exposure is to a supplier changing a model, its pricing or its terms, and to customer data being processed in ways the insurer did not agree to.

Fraud against the firm

The tools that help fraudsters fake claims evidence can also be turned on an insurer's own staff. In early 2024 an employee of the engineering firm Arup in Hong Kong transferred HK\$200 million, around US\$25.6 million, after a video call in which the chief financial officer and other colleagues were deepfakes. Claims payment, premium refund and supplier payment processes that rely on a voice or video confirmation need a second check that cannot be faked the same way.

People and skills

The World Economic Forum's Future of Jobs Report 2025 found that 41% of employers planned to reduce headcount where AI automates tasks. In insurance, Allianz Partners confirmed in July 2026 between 1,500 and 1,800 job cuts across Europe with AI cited as a factor, and in KPMG's September 2026 survey 36% of insurance leaders expected AI to reduce roles in claims by 2029 and 33% in policy servicing. Skills are short in the other direction: Capgemini found 67% of insurers citing a shortage of AI skills, and KPMG found 8% describing their workforce as highly proficient in AI. Evident recorded AI specialist roles at the largest insurers growing by 32% in a year while their total workforce shrank by 2.2%.



Capgemini, 2026 · KPMG, 2026 · Evident, 2026

Where the controls sit

Most of these risks are reduced by the same small set of controls. A person with authority approves decisions that commit money or determine cover. Every AI output is logged with the model, the data and the approver. Models are grounded in the firm's own policy and claims records and tested against a stated fairness measure. Evidence submitted by customers is checked for manipulation, and payment instructions are confirmed through a channel that cannot be faked by the same means. Almost all of those controls live in the firm's core systems and processes, which is why Chapter 6 treats the core system as the foundation of an AI programme and Chapter 9 builds on it.

Risks and the controls that address them

RISK	RECENT EXAMPLE	CONTROL
Bias and proxy discrimination	Citizens Advice ethnicity penalty research, 2022 and 2023	Test against a stated fairness measure before launch and monitor afterwards
Decisions the firm cannot explain	Air Canada chatbot ruling, 2024	Log every output with the model, the data and the approver
Outputs that are wrong	Deloitte Australia report refund, 2025	Ground models in the firm's own records and check output before it reaches the customer
Concentration and third parties	Bank of England review, 2025	Know where models run, how they change and how to exit
Fraud against the firm	Arup deepfake video call, 2024	Confirm payments through a channel that cannot be faked the same way
Decisions that commit money or cover	Every production example in this eBook	A person with authority approves them

Building an AI programme that delivers

For an insurer or MGA writing £20 million to £250 million of premium, the evidence in this eBook points to a programme built on a few well-chosen processes, clean data at source, AI bought inside platforms more often than built and governance proportionate to the decisions involved. BCG's finding that 70% of the difficulty in scaling AI lies in people and process, and Capgemini's that 72% of AI budgets go on technology, describe the same mistake from two directions.

Start with a process and a number

The programmes with published results each started with one process and one measure: triage minutes per submission at Zurich, settlement time on food spoilage claims at Allianz, liability assessment days on complex claims at Aviva. A firm that cannot state the current figure for the process it wants to improve cannot show that AI improved it. Measuring the baseline before any tool is chosen usually takes a few weeks.

The best first candidates in commercial lines are high-volume, document-heavy tasks where a person still makes the decision: submission intake and data extraction, bordereaux validation, first notification of loss, document generation and broker servicing queries. Each has a clear measure, a low regulatory profile and a short payback.

*A firm that cannot state the current figure for
a process cannot show that AI improved it.*

Fix the data where it is created

Models inherit the quality of the data they are given. Cleaning data downstream, in a warehouse or a spreadsheet before each model run, is slower and less reliable than capturing it correctly at the point of entry with validated fields. For an MGA that means structured submission and binding data in its own systems and agreed data standards with its coverholders and capacity providers; for a capacity provider it means requiring structured bordereaux in place of free-form spreadsheets.

Buy inside platforms before building

A firm of this size will not employ the data science teams that large insurers use, and does not need to. Most of the AI it will use will come inside the platforms it already runs or buys: the policy administration system, the claims system, the document and communications tools and specialist services for extraction, fraud and pricing. MIT NANDA's research, for all its limitations, found that partnerships reached deployment about twice as often as internal builds. The selection questions therefore move from features to architecture: whether the platform exposes its data through APIs, whether it records which model made which decision and whether the business can change the rules around a model through configuration.


Keep people at the points that matter

Outside simple, low-value personal lines claims, every production example in this eBook keeps a person at the decisions that commit money or determine cover. The rule is a sound default for regulatory reasons as well as operational ones: Articles 22A to 22D of the UK GDPR and POPIA section 71 both give customers rights where significant decisions are made solely by automated means, and a person with genuine authority to change the outcome takes the decision outside the strictest form of those rules. The approval step has to be real. A reviewer approving several hundred AI recommendations an hour is not exercising judgement, and a regulator or court will see that.

Govern in proportion

Governance for a firm of this size can be light, provided it covers five things. The firm should keep an inventory of every AI system in use, including those inside third-party software. Each system should have a named owner, who under SM&CR will be a senior manager for anything material. Systems that affect customers should be tested for bias against a stated fairness measure before launch and monitored afterwards. Outcomes should be tracked against Consumer Duty measures such as claims acceptance rates, complaint volumes and time to settle, broken down by customer group where the data allows. Suppliers should be asked where their models run, whether customer data is used to train them, what audit logs they keep and how the firm would exit.

Proportionate governance in five parts



Invest in the people who will use it

Capgemini found that 47% of employees given AI tools said their working day was unchanged after 18 months. Tools that are deployed without changes to the process, the role and the measures of performance tend to sit unused. Underwriters, claims handlers and operations staff need training in what the tools do and where they fail, and the processes they work in need redesigning so that the time the tools save goes somewhere useful. That work is the 70% in BCG's analysis, and it is where a mid-sized firm, with shorter lines of communication than a large insurer, has an advantage.

Use cases rated

The table rates the main use cases covered in this eBook on three measures, based on the evidence in Chapters 2 to 5. Evidence strength reflects how many named insurers have published measured results. Regulatory exposure reflects how directly the use case affects a customer's cover, price or claim. Time to payback is an indicative range for a mid-sized firm buying the capability inside an existing platform or service.

USE CASE	EVIDENCE	REGULATORY EXPOSURE	TIME TO PAYBACK	COVERED IN
Submission intake and data extraction	 Strong	 Low	Under 12 months	Chapter 2
Claims fraud detection	 Strong	 Medium	Under 12 months	Chapter 3
First notification of loss and claims triage	 Strong	 Medium	Under 12 months	Chapter 3
Underwriting guidance retrieval	 Moderate	 Low	Under 12 months	Chapter 2
Document generation and broker servicing	 Moderate	 Low	Under 12 months	Chapter 4
Bordereaux validation	 Limited	 Low	12 to 24 months	Chapter 6
Claims settlement on low-value, simple claims	 Moderate	 High	12 to 24 months	Chapters 3 and 5
Pricing models in commercial lines	 Limited	 High	Over 24 months	Chapter 2
Autonomous underwriting decisions	 Limited	 High	Over 24 months	Chapters 2 and 5

Readiness checklist

A firm that can tick most of these is ready to move a first AI use case into production.

Process and data

- ☐ We have chosen one process and recorded its current performance on a measure we already track.
- ☐ The data that process depends on is captured in structured, validated fields at the point of entry.
- ☐ Our core systems can pass data to an external service and receive results back through an API.

Ownership and oversight

- ☐ We keep an inventory of every AI system in use, including those inside third-party software.
- ☐ Each AI system has a named owner, and material ones sit with a senior manager under SM&CR.
- ☐ We know which decisions in the process commit money or determine cover, and a person with authority to change the outcome makes them.
- ☐ Customers affected by an automated decision can find out, make representations and ask for a person to review it.

Outcomes, people and evidence

- ☐ We have chosen a fairness measure and will test against it before launch and monitor it afterwards.
- ☐ We track Consumer Duty outcomes such as claims acceptance, complaints and time to settle for the process.
- ☐ The people who will use the tool have been trained in what it does and where it fails, and their process has been redesigned around it.
- ☐ Any image or document accepted as evidence in the process is checked for manipulation.

Questions to ask any AI or platform supplier

- 1 Where are your models hosted, and in which jurisdiction is our data processed and stored?
- 2 Is our data, or our customers' data, used to train or improve models used by other clients?
- 3 Which third-party model providers do you depend on, and what happens to our service if one of them changes its model, pricing or terms?
- 4 What does your system log for each AI output: the model and version, the input data, the output and who approved it? How long are logs kept and can we export them?
- 5 Can we explain an individual decision to a customer in plain language, and what information do you provide to help us do it?
- 6 How have you tested the model for bias, against which measure, and will you share the results?
- 7 How do you notify us before a model is changed or retrained, and can we test the new version before it goes live?
- 8 Can our business users change thresholds, referral rules and approval steps around the model through configuration, or does that need your development team?
- 9 How is the AI output written back into our policy, claims or billing record?
- 10 If we leave, what data and logs do we take with us, in what format and at what cost?

How Genasys approaches AI

Genasys is a cloud policy administration platform for insurers, MGAs, brokers and third-party administrators, headquartered in London with offices in South Africa. Chapter 6 set out five properties a core system needs if AI is to be used safely and at scale. This section describes how the Genasys platform handles each of them.

**450+**

DOCUMENTED RESTFUL API ENDPOINTS

**350+**

PRE-CONFIGURED PRODUCT TEMPLATES

**ISO 27001**

ACCREDITED AZURE INFRASTRUCTURE

Structured data

Policy administration, claims and billing share the same customer, policy and financial data, so a model reads one consistent record. Products are built from structured definitions using no-code tools and more than 350 pre-configured templates, and commission, claims and insurer bordereaux are produced from the same policy data, with outputs configurable to Lloyd's Coverholder Reporting Standards.

Open APIs and events

The platform has more than 450 documented RESTful API endpoints, published as OpenAPI with an interactive Swagger interface, behind a governed API gateway. Real-time event notifications and webhooks let an external AI service be called at the right step in a process and its output written back to the record.

A complete audit trail

Every policy transaction is logged, and in claims the audit trail records every action, approval and payment, which gives a firm the record it needs to explain and revisit a decision.

Configurable rules

Products, rating and rules are configured without a development queue, with version control so that every change is tracked and reversible. A referral threshold, an approval step or a new data field around an AI service is a configuration change.

AI kept separate from the core

Genasys keeps AI services separate from the core platform. It connects to AI-enabled services as each operation needs them and does not tie a client to one supplier; in claims, for example, data passes automatically to the fraud detection services the client chooses. As better models appear, they can be adopted without changing the core platform.

The platform is hosted in Microsoft Azure regions in the UK and South Africa on ISO 27001 accredited infrastructure, so data can stay in the jurisdiction its regulator expects.

SEE IT WORKING

Arrange a demonstration

To find out more about how Genasys connects the AI services you choose to your policy, claims and billing data, visit our website.

genasys.tech.com

Frequently asked questions

● How many insurers are using AI?

In the UK, 95% of the insurers in the Bank of England and FCA's 2024 survey used AI, though only 19 insurers took part. In South Africa, the FSCA and Prudential Authority found 8% of insurers using it, from a much broader base of firms. Internationally, BCG found 7% of insurers had scaled AI and about two thirds were still piloting.

● Where does AI deliver the clearest results in insurance?

In claims handling, fraud detection and submission intake. Published examples include an 80% cut in processing and settlement time on small food spoilage claims at Allianz, a 23-day reduction in complex liability assessment at Aviva and an 80% reduction in manual triage time on commercial submissions at Zurich.

● Will AI replace underwriters and claims handlers?

It is reducing some roles, particularly in high-volume servicing and claims work. Allianz Partners has cited AI in up to 1,800 job cuts, and 36% of insurance leaders in KPMG's 2026 survey expected fewer claims roles by 2029. Outside simple personal lines claims, the production examples in this eBook keep a person at the decisions that commit money or determine cover.

● Are there AI-specific rules for insurers in the UK?

No. The FCA has said it does not plan to introduce them. AI use is governed by the Consumer Duty, SM&CR and data protection law, including the new Articles 22A to 22D of the UK GDPR on automated decisions, in force since 5 February 2026.

● What about South Africa?

The FSCA has said an AI framework is being developed but had issued no guidance by October 2026.

Section 71 of POPIA restricts decisions with legal consequences made solely by automated processing, and the revised National AI Policy is expected for public comment in January 2027.

● Does the EU AI Act affect UK insurers?

Only where a UK insurer uses AI to assess risk or price life or health cover for individuals in the EU. Those obligations now apply from 2 December 2027. Commercial property and casualty lines are outside the high-risk category.

● What is agentic AI and should we use it?

Agentic AI carries out a sequence of steps towards a goal using tools and data. It works in high-volume, low-value, well-defined processes with a person approving anything that commits money or cover. Gartner expects more than 40% of agentic projects to be cancelled by the end of 2027.

- **What should a mid-sized insurer or MGA do first?**

Pick one high-volume process, measure it, make sure the data it depends on is captured in structured form at source and use AI that comes inside platforms with open APIs and a full audit trail. Chapter 9 sets out the approach in more detail.

References

All sources were accessed in October 2026. Entries marked (reported) were read through the named publication because the primary document could not be opened; entries marked (supplier account) come from a technology supplier's description of a client project.

Accenture, Why underwriters don't underwrite much, February 2022

Accenture, Underwriting Rewritten, August 2025

Accounting Times, Deloitte to refund government after using AI in \$440k report, October 2025

Allianz SE, When the storm clears, so should the claim queue, November 2025

Allianz UK, Allianz prevents record level of fraud, March 2026

American Bar Association, BC tribunal confirms companies remain liable for information provided by AI chatbot, February 2024

Association of British Insurers, Fraudulent insurance claims continue to top £1 billion, November 2025

Association of British Insurers, 2025 detected fraud figures (reported), Insurance Business UK, October 2026

Association of British Insurers, detected fraud series used in Chapter 3: 2015, 2016, 2017, 2018, 2019, 2020, 2021, 2022, 2023 restated

Aviva, Full year 2025 results announcement, March 2026

Aviva, Aviva stops record levels of claims fraud as scams become more sophisticated, June 2026

Bain & Company, A \$100 billion opportunity for generative AI in P&C claims handling, October 2024

Baker McKenzie, South Africa: regulators assess AI in finance, December 2025

Bank of England and FCA, Artificial intelligence in UK financial services 2024, November 2024

Bank of England, Financial Stability in Focus: Artificial intelligence in the financial system, April 2025

Bank of England, Response to the Treasury Committee inquiry report on AI in financial services, April 2026

Bank of England, SS1/23 Model risk management principles for banks, reissued April 2026

BCG, Insurance Leads in AI Adoption. Now It's Time to Scale., September 2025

BCG, AI-first companies win the future: property and casualty insurance, March 2026

Business Day, OUTsurance boss says human touch beats AI, September 2025

Business Insurance, Brokers tap AI to chase growth as rates fall, June 2026

BusinessTech, The new insurance scam that can get you blacklisted in South Africa, June 2025

Capgemini, World Property and Casualty Insurance Report 2024, April 2024

Capgemini, World Property and Casualty Insurance Report 2026, May 2026

Citizens Advice, Citizens Advice sounds the alarm on £280 car insurance ethnicity penalty, March 2022

Citizens Advice, People of colour still paying £250 more on car insurance per year, March 2023

Claims Magazine, BRIAN, the AI brain giving Allianz's commercial underwriters a head start, September 2025

Clyde & Co, A sector in motion: insights from the MGA market, July 2025

CNN, Arup revealed as victim of \$25 million deepfake scam, May 2024

Colorado Division of Insurance, SB21-169: Protecting consumers from unfair discrimination in insurance practices

Colorado General Assembly, SB 26-189, 2026

Computing, How AXA is using AI to rewrite the insurance model, 2026

Creamer Media, Naked raises R700m, January 2025

EIOPA, Opinion on AI governance and risk management, August 2025

EIOPA, Survey on generative AI shows swift but cautious adoption among Europe's insurers, February 2026

Entrepreneur (reporting Bloomberg), Klarna CEO reverses course by hiring more humans, May 2025

European Union, Regulation (EU) 2024/1689 (AI Act), July 2024

European Union, Regulation (EU) 2026/1744 (Digital Omnibus on AI), July 2026

Evident, AI Index for Insurance 2026, June 2026

EY-Parthenon, GenAI in Insurance: Spring 2025 survey, 2025

FAnews, Protecting your clients in the age of generative AI and deepfakes, December 2024

FAnews, The R3.5 billion challenge facing SA insurers and intermediaries, January 2025

FAnews, Insurers, intermediaries and UMAs stride into the AI age, June 2025

FAnews, Using AI to fight fraud in claims, November 2025

FAnews, How collaboration and shared intelligence is crushing insurance crime, June 2026

FCA, Our approach to AI, updated October 2026

FCA, AI Lab, updated July 2026

FCA, Supercharging the digital sandbox, June 2025

FCA, FCA helps firms test AI safely, December 2025

FCA, FCA announces second cohort of AI Live Testing, April 2026

FCA, Consumer support outcome: good practices and areas for improvement, March 2025

FCA, Research note: literature review on bias in supervised machine learning, December 2024

FCA, The FCA's long term review into AI and retail financial services, speech by Sheldon Mills, January 2026

FCA, Insurance Regulatory Priorities, February 2026

FCA, The Mills Review: AI and the future of retail financial services, July 2026

FCA and Yonder, AI consumer research: quantitative, 2026

FSCA and Prudential Authority, Artificial Intelligence in the South African Financial Sector, November 2025

Gallagher Re, Global InsurTech Report Q1 2026, 2026

Gallagher Re, Global InsurTech Report Q2 2026, August 2026

Gartner, 64% of customers would prefer that companies didn't use AI for customer service, July 2024

Gartner, Lack of AI-ready data puts AI projects at risk, February 2025

Gartner, 50% of organisations will abandon plans to reduce customer service workforce due to AI, June 2025

Gartner, Over 40% of agentic AI projects will be cancelled by end of 2027, June 2025

Genasys, genasystech.com: product, connectivity, policy administration and claims management pages, October 2026

HM Government, AI regulation: a pro-innovation approach, March 2023

Hiscox, Hiscox and Google Cloud AI lead underwriting, December 2023

Hiscox, AI underwriting model live, August 2024

Hiscox, Full year 2025 results, February 2026

HSBC, Beyond the bubble, December 2025

ICO, Consultation on draft guidance on automated decision-making, March 2026

Insurance Age, Only 20% of brokers working with AI, May 2025

Insurance Age, IFB warns on sharp rise in ghost brokers, May 2025

Insurance Business Australia, Hollard tests AI to streamline consultants' claims file reviews, March 2026

Insurance Business, MGAA honours five MGAs as FCA oversight review puts delegated authority under the microscope, July 2026

Insurance Business, Inside Marsh's AI build, August 2026

Insurance Business, WTW bets US\$625 million on AI, July 2026

Insurance Business, Allianz confirms job cuts as AI reshapes insurance, July 2026

Insurance Business, London market's CDR now covers full insurance lifecycle, October 2026

Insurance Edge, Davies deploys AI-powered claims intake agent, April 2026

Insurance Post, Esure on generative AI in the contact centre, September 2024

Insurance Times, Markel reduces quote turnaround time with help of AI, September 2023

Insurance Times, Allianz to enhance fraud prevention as scammers turn to AI and deepfakes, 2024

Insurance Times, Unpacking the ethnicity penalty in car insurance premiums, 2023

IOL, SA insurers expose fraud, July 2025

IOL, How rising insurance fraud is reshaping the industry in South Africa, May 2025

ITWeb, Insurers deploy AI to tackle fraud, July 2026

IVANS (article by Cytora), Arch Insurance AI submission intake (supplier account), May 2026

Kalepa, Paragon doubled its quote-to-bind rate (supplier account), March 2026

KPMG, Unlocking AI value in insurance, September 2026

KPMG and University of Melbourne, Trust, attitudes and use of AI: South Africa snapshot, 2025

Lemonade, Form 10-K for 2025, 2026

LIMOSS, Delegated Data Manager

LMA, Operations Q1 2026 update, April 2026

LMA and Barnett Waddingham, AI adoption more than doubles across the Lloyd's market in 12 months, April 2026

McKinsey, Aviva: rewiring the insurance claims journey with AI (adviser's account)

McKinsey, The state of AI in 2026, August 2026

Microsoft, Gallagher customer story (supplier account), April 2026

Microsoft, Markerstudy Group embraces AI (supplier account), November 2024

MIT NANDA, The GenAI Divide: State of AI in Business 2025, July 2025

Moonstone (reporting the FSCA and Prudential Authority survey), Regulators step in as AI adoption exposes new financial risks, April 2026

NAIC, Model Bulletin adoption map, October 2026

OpenAI, Travelers customer story (supplier account), June 2026

Reinsurance News, Zurich scales Cytora AI platform across global underwriting operations (supplier account), May 2026

Reinsurance News, Insurance sector faces major gap between AI confidence and transformation: KPMG, September 2026

SA Government News Agency, Minister announces withdrawal of draft AI policy, April 2026

Sanlam, Sanlam gives AI a seat, March 2026

Santam, 2025 results presentation, 2026

SAS, Synthetic images and AI insurance fraud, May 2026

Tech.eu, Klarna boss on the power of AI, February 2024

UK Parliament, Data (Use and Access) Act 2025, section 80, 2025

UK Government, The Data (Use and Access) Act 2025 (Commencement No. 6) Regulations 2026, January 2026

World Economic Forum, Future of Jobs Report 2025, January 2025

Zawya, Naked launches ChatGPT app for car insurance quotes, 2026

Zurich UK, Zurich detected £260,000 worth of bogus insurance claims every day, March 2025

Where AI is delivering in insurance, and what it takes to use it well.

Most insurers already use AI somewhere in the business. Few have scaled it, and fewer still can show what it returns. This eBook sets out where AI is delivering in 2026, what it takes to use it well and what that means for firms writing £20 million to £250 million of premium in the UK and South Africa.

It covers underwriting, claims and fraud, distribution and agentic AI, with six case studies that report measured results. It then turns to the data and core systems AI depends on, the rules that govern it in the UK, the EU, South Africa and the US, and the risks that need managing.

It closes with a practical approach for a mid-sized firm: a rating of the main use cases, a readiness checklist and ten questions to put to any AI or platform supplier.

genasystech.com

LONDON | JOHANNESBURG | CAPE TOWN