

INFOSYS GLOBAL
INSURANCE JOURNAL
VOLUME 2



Table of contents

Foreword	4
Executive summary	6
AI in insurance: Tried and tested, but rarely scaled	10
How AI will drive transformation for insurers	20
How to capture data and knowledge for humans and agents	28
Ways to govern AI risk in insurance	36
Buy, price, settle: The new arithmetic of AI ROI in insurance	44



Kannan Amaresh
Senior vice president and
head of global insurance, Infosys

Foreword

Technology has always advanced, but rarely this fast. AI is being rolled out across industries as leaders race to shift from [experimentation to creating value](#) — a transition that once took decades is now measured in quarters.

For the insurance industry, the challenge is particular. The industry is built on long-established processes, risk controls, actuarial judgment, and regulatory trust. Yet AI adoption has outpaced that of personal computers, the internet, and smartphones combined. Integrating the fast new technology into this seasoned, structured industry requires a deep understanding of both.

As you will read in our second Infosys Insurance Journal, insurers are using AI now, and looking for ways to scale it, improve

efficiency, and increase productivity. Generative AI tools have the potential to offer new solutions to fast-changing challenges that insurers already face. These challenges include climate change, demographic shifts, and geopolitical uncertainty.

As is true with all transformational technologies, existing industries must retool and alter incumbent processes to achieve real value. Insurers and software specialists are finding AI valuable for digital modernization, but this effort must include worker reskilling and retraining. The goal is to merge the best longstanding practices with the proper use of modern technology.

AI, of course, must be integrated responsibly and comply with existing regulations. This must be done at a time when new rules and requirements for AI remain in flux.

And finally, insurance executives and their boards expect great things from AI very soon. Insurers must develop ways to measure their return on AI investment that keeps insurance core functions at the heart of decision-making.

The future of AI in insurance will be driven by talent and token economics, where scarce high-quality talent defines how to make AI effective, and token economics determines

how to keep it economic. Together, these dual forces will shape the pace of innovation and the sustainability of AI adoption.

AI may yet buck the trend of slow transformation as enterprises race to implement it. But measurable value still requires time and hard work. Insurance executives must begin now to convert early AI use into measurable, responsible productivity gains.





Executive summary

Insurance enterprises take their data, technology, and risks seriously. This has long been reflected in the industry's measured approaches to data management, technological change, and risk assessment.

Savvy insurers now recognize that the entrance of artificial intelligence (AI) into their industry dictates a shift from the measured approach of the past.

AI is proliferating across the insurance value chain and showing potential for great transformation. Which brings us to the current test of insurance and AI: Its use in the industry is widespread, but its outcomes are as uneven as AI's own [jagged frontier](#). This journal examines where the insurance industry stands with AI, the possible future, and what it will take to achieve the greatest potential value.

Tested, tried, and rarely scaled

Insurers have adopted AI more eagerly than any other emerging business technology of the 21st century. Underwriters use AI to summarize documents and bring structure to unorganized information. Claims teams automate routine tasks. IT departments are using it to modernize legacy code and bust silos. Yet these uses have not coalesced into enterprise-scale transformation. [Pilot fatigue](#) has been a documented problem since 2025. The cycle of pilots that succeed and fail to reach production repeats and repeats.

Legacy systems, fragmented data, governance shortfalls, lack of training and poor implementation — all share some blame. More significantly, AI budgets grow larger, return expectations draw nearer, and the question over whether to modernize

something AI could render obsolete looms larger.

There is hope. Enterprises are beginning to realize that **organizational transformation** must accompany AI adoption to achieve scale. This transformation must span the enterprise and include combined infrastructure, agreed governance, workforce reskilling and clear business goals.

New tech for new risks

Climate volatility, demographic change, heightened cyber threats and geopolitical instability will make risk harder to price. More assets and exposures will be deemed uninsurable under conventional approaches. AI, in concert with other emerging technologies, will give insurers ways to adapt. These range from real-time pricing, and hyperpersonalized coverage, and parametric products that pay out when predetermined conditions are met.

The insurance workforce will be transformed by AI as well. AI agents — with appropriate human supervision — will take on routine tasks and free humans for higher-value activities. This only works if insurers keep trust at the core of their activities. Every action involving AI must come with explainability and transparency for insurers to maintain their reputations, the confidence of their customers, and compliance with regulations.

Data capture and wisdom

Insurers can only realize the future-horizon possibilities of AI if they get their data right

and preserve the wisdom of their present-day professionals and experts. At present, this is a real risk. Decades of merger-driven growth have left companies with unstructured data and incompatible legacy systems. Institutional knowledge comes and goes every day in the minds of experienced employees, many of whom are nearing retirement. Capturing data and wisdom means building canonical data models that give AI agents a single consistent source of truth. It means capturing tacit knowledge and veteran know-how before it walks out the door forever. And it means building processes and governance guardrails for humans and AI agents to collaborate safely and confidently. These models and processes will capture and enhance the data and wisdom that set top-performing insurers apart.

How to govern AI in insurance

Along with the potential advantages AI brings to insurance, it also introduces new hazards. This begins with the fundamental nature of large language models and generative AI. Their probabilistic nature runs against the grain of conventional deterministic insurance practices. Insurers and their regulators are grappling with understanding this and the second- and third-order implications of AI and AI agents. At this stage, regulatory frameworks and required guardrails vary widely by jurisdiction. For example, under the European Union's AI Act, AI systems for life and health insurance pricing have been deemed high risk and must operate with human oversight, documentation and other compliance requirements. In the US, federal policy aims to reduce AI regulatory

burdens, but at times conflict with state laws and guidance from industry associations. The United Kingdom is working toward a principles-based and technology-agnostic approach to AI and financial services, and Singapore has established a relatively permissive but structured framework for fairness and explainability. In all cases, using AI does not exempt insurers from existing rules and requirements.

The new ROI equation

Insurers have adopted AI quickly. The key question is quickly shifting from “does it work?” to “is it worth it?” Those who figure out the math for return on AI investment will separate themselves from the pack.

Boards and executives expect AI investments to pay off quickly. But the constants and variables are not yet fully defined. Insurers can better calculate AI value by grounding the discussion in three key insurance domains: How customers buy policies, how insurers price risk, and how they settle claims. This gives insurers familiar territory to measure and prove where and how AI pays off.

Enterprise AI has arrived and offers opportunities for insurers to operate more quickly, improve customer experience, and be more productive. New frameworks for effective AI use now need to be built. This requires new operating models, digital transformation, and workforce retraining. Still, time-tested frameworks and actuarial truths must be retained. Above all, insurers need to preserve trust as their core principle and put AI to work in responsible ways.





AI in insurance: Tried and tested, but rarely scaled

- AI use is now widespread across insurance, with many companies deploying or piloting it across business units.
- Despite significant investment, large-scale adoption remains limited as insurers struggle to move beyond pilots, held back by legacy systems, fragmented data, governance challenges, and skills shortages.
- Insurers achieving the greatest success treat AI as an enterprise transformation program, combining strong governance, shared infrastructure, workforce engagement, and clear business objectives to help scale it across the organization.

In just about every area within insurance, companies are investigating how to use artificial intelligence (AI). Underwriters are using generative AI to create document summaries; claims departments are piloting automation projects and using AI to stop fraud; IT teams are using it to help modernize legacy systems; actuaries are using it in pricing; and business intelligence units are using AI models to improve reporting and analytics.

The speed at which insurers are being drawn to this new technology is unusual for the sector. In our experience, insurance has been a noted laggard when it comes to adopting other innovations. But this should not surprise. As a sector, insurance generates large amounts of structured and unstructured data, which is perfect to build AI models. AI is also useful in supporting decision-making. In addition, insurance has been able to look to other sectors such as banking for efficient

AI use cases, including fraud detection, risk analysis, and knowledge-based functions, such as underwriting, actuarial analysis and claims adjudication.

Based on Infosys's experience working with insurers, companies expect AI to create value, reduce costs, increase efficiency, and improve performance. These are needed urgently to address [mounting pressure](#) from rising claims costs, skills shortages, and competition from digital-native insurers such as US-based [Lemonade](#), which typically operate more efficiently. Further, customers are demanding faster, more personalized, and more transparent services similar to those available in retail and banking. Insurers' efforts to generate value with AI have been partial. Some 17% of insurance respondents said agentic AI delivered significant positive impact for them in terms of cost savings, according to Infosys Knowledge Institute research (Figure 1).

A long shadow


Despite the evident draw, [questions over](#)

[returns on investment \(ROI\)](#) are casting a shadow over full adoption for some insurers. Even where ROI has been estimated, there's no guarantee it will be achieved. For example, one estimate is that generative AI use alone could [unlock \\$50 billion to \\$70 billion of new revenue](#) for the sector, with the highest impact on marketing and sales, customer operations, and software engineering. But gains like these depend upon [an enterprisewide AI strategy](#).

Most insurers have not advanced from standalone projects to enterprisewide transformation (Figure 2). [Few of their AI pilots and projects are creating value](#) — increasing the top line of sales. Instead, many AI initiatives are aimed at cost-cutting and activities such as customer service chatbots and document summarization. These are useful, but far from transformative.

In our experience, this introduces insurers to new challenges such as duplicated models, inconsistent governance frameworks, fragmented data pipelines, and increased maintenance costs. It can also lead to greater

Figure 1. Agentic AI impact on cost




N = 35

Source: Infosys Knowledge Institute

Percentages do not add to 100% due to rounding.

Figure 2. AI use in insurance tends toward point solutions and productivity gains



N = 187 use cases

Source: Data from Evident AI Index Insurance, June 2026

unknown liabilities, as AI systems embed flawed assumptions, unexplained decisions, and misaligned incentives into infrastructure or processes. Such risks increase when AI initiatives are developed in departmental silos instead of being governed centrally.

Insurers are also [experiencing pilot fatigue](#) from running pilots that fail to scale, even when they are technically successful. Hurdles include ineffective governance, compliance requirements, procurement constraints, and systems integration challenges. When this occurs, budgets are drained and staff become frustrated, making it successively harder to implement enterprisewide change.

Nevertheless, AI budgets have been [rising rapidly](#). Yet it remains unclear whether this represents genuinely new investment or a redistribution of existing technology spend away from core infrastructure, modernization, and maintenance. Infosys executives describe insurers debating whether they continue spending large proportions of IT budgets on

sustaining legacy systems or shift those funds to AI spending. There is a risk that aggressive AI spending could create further technical fragmentation if it is not integrated into broader enterprise transformation programs, Infosys executives say.

Linked to these challenges are questions about the value for insurers of modernizing their platforms now, just as advancing AI capabilities might make some capabilities obsolete in the not-too-distant future. This uncertainty could be putting a brake on investment, as insurers seek the most cost-effective balance between revamping systems and creating value from that spend.

Other challenges insurers face include overcoming expensive-to-maintain, fragmented, and undocumented legacy systems that are no longer supported by the vendor, particularly where companies have grown through acquisition, poor data quality, workforce resistance, and AI skills shortages. In addition, insurers face governance and

compliance issues with consistent [concerns about explainability](#), fairness, security, and accountability.

AI to modernize

Some insurers are using AI to play catch-up and bring down costs. [This includes using generative AI to analyze code](#) and generate structured documentation, as well as helping companies preserve institutional knowledge, which is increasingly essential as the workforce ages and/or leaves. Others are using generative AI to automate code creation and testing, while agentic AI is being used to automate complex workflows, Infosys insurance experts say.

Reported benefits include one insurer applying generative AI to improve coding and testing efficiency, [halving the time taken](#).

Yet modernization itself is becoming more strategically uncertain. Insurers are [investing heavily](#) in replacing legacy infrastructure with cloud-native and AI-enabled systems, but

rapid advances in generative and agentic AI raise questions about how durable those investments will be.

While few insurers are delaying modernization programs, the pace of AI development has introduced uncertainty over what future-ready infrastructure should look like. [Some](#) are therefore adopting a phased modernization strategy, combining cloud migration with modular AI deployment, rather than attempting full system replacement. This recognizes that AI transformation is as much an organizational challenge as a technical one.

How to combat pilot fatigue

There are some standout examples of insurers taking an enterprise approach to AI within their operations.

Zurich Insurance has adopted a group-wide AI360 strategy to scale AI across underwriting, claims, customer service, and operations to avoid the pilot trap. Today it has more than [460 AI use cases](#) in production or deployed.

While few insurers are delaying modernization programs, the pace of AI development has introduced uncertainty over what future-ready infrastructure should look like.



Key elements to its strategy are:

- A formal enterprise-governance framework known as STAR — safety, transparency, accountability, and reliability.
- AI literacy programs for staff — “AI for everyone,” including a £1.3 million (\$1.7 million) investment into an AI apprenticeship program in partnership with Imperial College London.
- Embedding AI into core workflows.
- Dedicated AI research partnerships with universities — Zurich AI Lab, with the University of St. Gallen and ETH Zurich. It also runs the Zurich Innovation Championship, a global startup program through which Zurich pilots and scales innovative solutions across its operations.
- Structured processes from developing and testing prototypes into adoption.

Including AI, Zurich’s digital-transformation spend has already exceeded \$1.8 billion, with broad, cross-cutting outcomes, including faster claims processing and higher customer satisfaction.

Zurich’s approach demonstrates that scaling AI depends less on individual tools and more on getting the governance right, engaging the workforce, and ensuring alignment across the whole organization. Similarly, its “AI for everyone” program, with its emphasis on getting the workforce to understand AI, acknowledges the importance of cultural

transformation when it comes to this new technology.

Silo elimination

Another company successfully adopting AI at scale is AXA. At its core is a [shared AI infrastructure](#) with reusable AI products, components that can be deployed and used across multiple business units, teams, or use cases. This includes an enterprise platform that 140,000 employees can access and a sourcing model to promote the use of AI across different business units.

The company also buys AI products. For example, its health agents historically spent up to six minutes after a call writing it up. This important but time-consuming task also increased the risk of record-keeping errors. The solution was to adopt an off-the-shelf bot from Verint, trained on AXA's own large language model using in-house data, to transcribe and summarize customer calls. This cut average handling time per call by 60 seconds and gave agents more time to spend on customer interactions rather than administrative tasks.

AI has been a core technology for Chinese insurance giant Ping An since 2017. The company uses machine learning and advanced analytics to assess [global climate risk](#), predict weather patterns, provide alerts, and expedite claims processing.

Ping An claims that this strategy has delivered:

- A life and health claim service with one-

sentence reporting, one-click uploading, and one-minute validation; policy underwriting within seconds for 94% of cases; and average claim processing times of 7.4 minutes.

- Vehicle damage video-assessment tools with 95% accuracy, reducing time to process claims from hours to minutes, improving customer satisfaction and cutting costs.

Good governance

Insurers need transparency to operate accountably, build trust with customers, and answer to regulators. Indeed, four of the biggest regulatory concerns about AI in insurance are [bias, explainability, customer outcomes, and transparency](#). But it's not just regulators who care about these things; customers can also challenge decisions. All this together means that [any use of AI by insurers must include explainability](#) — how any outcome was achieved.

As an indicator of the importance of transparency, intelligence platform Evident Insights set up the [Evident AI Index for Insurance](#). This ranks insurers as they build and deploy AI, assessing key capability areas of talent, innovation, leadership, and transparency, alongside maturity. Allianz scores well across all four areas, second only to AXA, reflecting its own clear focus on organizational and governance issues as it ramps up AI use.

Helping to boost its score, Allianz has [eight core principles](#) for responsible AI

use that include transparency as one and accountability with accuracy and proficiency as another.

Under transparency, it commits to being open about the purpose of its AI systems and to tell customers when they are interacting directly with an AI system such as a chatbot. When it comes to accountability, Allianz monitors outputs of its AI systems to ensure they match the intended outcomes, including integrating user feedback.

Such an approach indicates a shift toward seeing governance and transparency not as a barrier to adopting AI, but as a prerequisite to scaling it up — and a marketing advantage.

Adoption resistance

When staff are concerned that a new technology will disrupt their jobs, they can be reluctant to embrace it. AIG has acknowledged this and taken an approach to using [AI that puts staff at the center](#), implementing solutions that enhance workers' roles.

For example, its generative AI-powered tool, developed with software company Palantir and AI company Anthropic, ingests and prioritizes excess and surplus claim submissions. In effect, it allows underwriters to review more policies without hiring extra staff, preserving jobs while increasing efficiency.

AI adoption across insurance is gathering pace, but scaling remains difficult. Leading

insurers are succeeding because they take a structured approach that combines governance, infrastructure, culture, and strategy.

How to scale AI

The insurers successfully deploying AI at scale tend to have a similar approach, underpinned by a strategy that sees AI as a catalyst for organizational transformation rather than simply part of ongoing technological change and improvements.

These successful practices include:

Focus on fewer, high-value enterprise-scale use cases. Instead of disconnected pilots, leading insurers focus on a few high-value use cases that deliver measurable operational or commercial results. In this way they reduce duplication and make governance easier to manage and monitor, while the scale of the projects creates momentum that encourages wider adoption.

Build shared AI platforms and data layers. They develop shared infrastructure that can cross between business units, avoiding duplication and waste. This improves consistency across models, makes governance simpler and easier, and aids data management. Centralized data layers and reusable AI tools are particularly important for insurers with fragmented systems arising from acquisitions or differing regional structures. This approach also allows insurers to monitor models more consistently and scale success faster.

Create formal, responsible AI frameworks.

When it comes to governance, success demands being on the front foot. Insurers at the forefront of using AI have developed governance frameworks that ensure explainability, fairness and bias, accountability, transparency, and security. They operate with an eye on compliance and the need to keep the regulators happy, while also managing reputational and operational risk. Additionally, protocols that can identify the optimal way to make output from a particular AI model more explainable, for example, are in [the early stages of development](#) and likely to prove useful in future.

Invest heavily in AI literacy and training.

Insurers successfully getting their staff to adopt AI take the time to explain how it will affect their jobs, the company, and the sector, and train them to use and work alongside AI systems with confidence. It's also helpful to educate those not directly affected and to be open and transparent about the timetable for AI's adoption, where and how it will be used, and the benefits it will bring across the organization.

Embed AI directly into workflows.

Rather than deploying AI in standalone applications, leading insurers are integrating them into existing operational processes — and redefining these as required. AI is also being embedded into workflows, which increases usability.

Prioritize measurable operational improvements.

Success is boosted when there are clearly defined goals and measurable outcomes for staff using AI. These goals might be reduced claims handling times, lower administration costs, improved fraud detection, or increased underwriting capacity. But they need to be easy to monitor and measure, and link to clear key performance indicators.

Centralize standards while enabling local deployment.

Insurers are reducing operational risk by developing enterprisewide standards for governance, infrastructure, and security, while taking account of regional or local regulations and product lines.

Centralized data layers and reusable AI tools are particularly important for insurers with fragmented systems arising from acquisitions or differing regional structures.

Develop long-term enterprise AI operations capabilities. Successful insurers are looking to the longer term when they are building capabilities to manage monitoring, governance, model updates, cybersecurity,

and integration across their organization. This helps create an enterprise culture of AI adoption rather than acceptance by individual teams — and sets the stage for successful adoption at scale.



How AI will drive transformation for insurers

- Trust is the bedrock of the insurance industry: AI models must be made more transparent and explainable to speed their adoption.
- AI will multiply human capability in a hybrid future workforce, increasingly made up of humans working with AI agents.


Artificial intelligence (AI) is set to enable two profound shifts in the insurance industry over the next two decades.

First, the nature of risk will transform, led by climate volatility, demographic changes, systemic cyber risks, and geopolitical instability. These will all make risk harder to price — and some assets uninsurable. Insured losses caused by natural catastrophes, for example, are rising between 5% and 7% a year on average and **could reach \$186 billion by 2030**, with the US and Europe both hotspots (Figure 1).

And second, there will be a shift from selling annual insurance policies to **continuous coverage, updated in real time as customer needs change**. This will accelerate as insurers offer new products that match more technology-enabled lives, such as allowing drivers to insure individual journeys when they use car-sharing services, or adjusting premiums based on driving routes chosen.

Rather than responding after the event, insurers will become increasingly proactive, supporting customers to prevent harm or making fixed payments without a claim

Figure 1. Insured losses due to natural catastrophes



Source: Data from [Swiss Re sigma 1/2025: Natural catastrophes: insured losses on trend to \\$145 billion in 2025](#)

needing to be made, through parametric products which are based on preagreed conditions such as a certain level of rainfall. The African Union's African Risk Capacity agency, for example, is already [using parametric payments for farmers](#) in countries where crops are ruined by drought.

Faced with change on such a scale, [AI is the technology](#) that will underpin these shifts. It will make possible more accurate and timely pricing of risk by accessing and processing larger volumes of real-time unstructured data. It will deliver hyperpersonalized insurance coverage, such as in life insurance, where health data collected over a number of years from wearable devices and other types of monitoring could be used to inform [more accurate and customized pricing and underwriting models](#). This will be increasingly important for companies wrangling with the challenges of insuring the cost of an aging population — between

2015 and 2050, the proportion of the world's population aged over 60 [will nearly double from 12% to 22%](#).

AI will also play an important role in enabling embedded insurance (offering insurance on a product at the time of purchase), and in assessing and resolving claims faster, with the aim of improving the experience for customers, reducing costs, and enabling new business lines.

Finally, both workflows and workforces will be transformed by the improvements in speed and accuracy that AI tools and agents are expected to bring.

However, change at such a scale inevitably comes with risks for insurers. A key priority will be maintaining trust as they automate decision-making, because reputation is paramount. Insurers must also prepare for [far more sophisticated cybersecurity risks](#).

Below we consider these challenges in detail, along with ways insurers can prepare for what's to come. This will allow the pace of AI adoption to safely accelerate in the years ahead, supporting the industry through its time of radical change.

New approaches to data

Insurance is an inherently data-driven industry: the entire business model depends upon the skill of actuaries. Better-quality data leads to more accurate risk assessment and fairer premiums. On the claims side, gathering the facts of an incident as quickly and conveniently as possible allows insurers to assess claims and process them in a timely way.

Further, insurers deal with vast quantities of unstructured data, such as emails, text messages, and documents. Traditional databases struggled to store and process data from sources such as these, but AI tools can now do it faster and more efficiently than [conventional methods](#). However, this technology is not yet being used at anywhere near its potential, so companies are missing out on savings and potential benefits. That is likely to change in the years ahead.

Using generative AI to its full potential in claims handling, for example, could reduce loss-adjusting expenses by between 20% and 25%, and leakage (the difference between what is paid and what is owed on the insurance contract) by between 30% and 50%, [creating more than \\$100 billion in benefits](#) for insurers and customers, according to one estimate.

Swiss Re's [ClaimsGenAI tool](#) already picks up documents immediately when they arrive in the company's claims department, scans them for useful data and keywords, and organizes the information to boost claims handlers' efficiency. It is built on insights from more than two decades of unstructured claims data, so the tool is also capable of detecting potential fraud or recovery opportunities that manual analysis might not have picked up. If similar technology is applied across all insurers, this could speed up claims handling.

Looking ahead five years, 10 years, and beyond, insurers will be able to gather data from Internet of Things (IoT) sources and sensors to offer continuous pricing for home and vehicle insurance. Policies might, for example, come to include features like active loss prevention, such as valves that can shut off water in a house when a leak is detected. This could reduce premiums as well as total losses to insurers.

Insurers will also gain expertise in pricing the risk of autonomous vehicles as they become part of day-to-day traffic. In Singapore, for example, Allianz has started on this path by [insuring the country's first fully autonomous bus](#). The policy will cover any damage caused to the bus's cameras and sensors, as well as any compensation that operator WeRide may need to pay for damage to property, personal injury, or death caused by the self-driving vehicle.

In health and life insurance, health monitoring to reduce premiums is likely to tip further into the mainstream, as the

challenge of insuring increasingly aging populations extends further into [middle- and low-income countries](#). Additionally, customers are likely to live for longer with a range of chronic conditions. According to one study, more than half the global population aged 60 and over [suffers from multiple chronic conditions](#).

There will also be a smaller working-age population of people paying insurance premiums to support the cost and consequences of this demographic change. To help overcome these challenges and improve pricing accuracy in underwriting processes, Willis Towers Watson, an insurance broker and advisory company, last year started working with the health-data analytics provider Klarity to create a tool that produces individual-level scores to predict and classify mortality risks. It enables real-time insights into individual health risks through [integrating data from wearables such as smart watches and other devices](#).

Explainable AI, insurance

AI has transformational potential for insurance. However, it also comes with significant risks for the industry, of which reputation is the greatest.

Insurance is a heavily regulated industry and many of the largest companies have built up their brand [over 100 years or more](#). As a result, caution abounds — one survey found that 87% of senior

insurance professionals across the UK and Europe are “concerned” about [bias or unfair outcomes in AI-driven processes](#). In the experience of Infosys experts, insurers are advancing into AI use more warily than retail, telecommunications, and financial services companies.

Insurers have moved forward slightly faster in adopting AI for their internal processes, but remain more cautious when it comes to customer touchpoints and compliance and regulatory risk functions. For these, they have preferred to maintain a high degree of human supervision and intervention. As companies look to scale AI adoption in the years ahead, putting insurance-specific governance guardrails in place is essential. These should ensure decisions made by AI agents over pricing or why a claim has been rejected are transparent and explainable to customers.

Allianz, which has more than [900 internal AI use cases](#) registered worldwide, has a partnership with Anthropic [to advance responsible AI](#). A key element is that as AI systems are built, the decisions the tools make, their rationale, and the data sources used will all be logged, so they are traceable and meet regulatory compliance standards. As more processes are handled by agentic AI, [transparency and explainability will only become more important](#). The ambition is to set new standards for accuracy that can establish a benchmark for both AI developers and insurers.

The company has also made public its [eight principles for responsible AI](#), which cover transparency, accountability, security, non-


discrimination, data privacy, data governance, human oversight, and applications where AI use is not allowed. These include “social scoring” of individuals seeking to buy insurance based on their inferred traits or behavior, assessing employees at work using AI emotion-recognition tools, and using AI to circumvent legal and regulatory safeguards.

Another notable risk for insurers now and in future is from increasingly sophisticated AI-driven cyberattacks. As AI’s capabilities evolve, the risk posed to insurers’ customer and corporate data by criminals using the technology is increasing.

Insurance’s future workforce

Workforce implications will be another consideration for the industry as insurers introduce AI in more areas across their businesses and deepen access. Job losses are inevitable, as parts of claims handling and policy administration are carried out by AI agents. Estimates are that insurers could make [net efficiency gains of 30% to 40% by using AI to change the structure of their workforce](#), with, for example, 50% fewer customer service roles, but a 20% increase in both product management and IT roles. Infosys [AI Business Value Radar](#) research in 2025 found that some of the most viable AI use cases are developing in the insurance industry. Further, the report showed that extensive levels of workforce preparation makes AI more effective (Figure 2). Despite the shift AI potentially brings to insurers’ operating model, humans are still very much required. An underwriting assistant will still work in underwriting, but they will also be an

Figure 2. Workforce prep makes AI more effective



N = 3,798

Source: Infosys Knowledge Institute AI Business Value Radar
Percentages reflect the probability that AI uses achieve most or all objectives at each tier.

AI reviewer, while a claims handler would also be an exception manager, and an actuary would also be a model validator. In these cases, AI will act to amplify the talent within an organization, moving humans to higher-value work and allowing staff to expand their roles and responsibilities.

The AI-powered future

Insurance leaders know they are facing a future marked by disruption, from elements as varied as climate events and demographic changes. The arrival of AI at scale across all sectors also brings pressure from customers to offer the interactions and rapid transactions they get from other businesses in more advanced sectors. Competition will intensify between insurers too, as early AI leaders gain benefits in productivity and offer better customer experience.

Looking ahead, insurers need to learn how to use AI and [leverage AI agents — autonomous software systems that can perceive, decide and act to achieve specific goals at scale](#). For

example, wherever an adverse climate event occurs, whether in North America or Africa, parametric insurance products leveraging AI will be available that can assess liability and make payments automatically. Across the board, claims processes will become simpler and more automated.

Prepare for the AI-first future

Companies can establish solid foundations now by focusing AI investments on the parts of their business where the technology will [bring the greatest benefits](#). To deploy AI in a structured and scalable way, companies should:

Consider how best to apply AI to keep pace with changing demands. AI will allow insurers to be proactive about potential claims, rather than responding after the event. It will also enable dynamic pricing and faster claims and settlement, or in specific cases, parametric payouts. AI can be used to automate some low-risk tasks such as renewal alerts, quote assistance, and claim status communication, while for more complex tasks, AI could be embedded to automate certain steps of fraud detection or underwriting risk assessment, for example. The most complex tasks will be executed by humans but augmented by AI to make processes faster in, for example, loss estimation and litigation claims.

However, first identify the need. In the case of high-risk programs with a technology capability that works, and which don't require many humans, will AI really make things so much better that it repays the investment?

Infosys' [AI Business Value Radar](#) found that one-fifth of AI uses returned no value or had been canceled before scaling. When AI does not deliver business value, there is limited point in insurers assessing any other kind of business benefit, such as staff or customer retention. Technology-led businesses such as Uber have already [capped individual employee spending on AI](#) as they seek to moderate the high cost of using it until the returns on investment are clearer. AI is the technology dominating the headlines, but it is not the only tool that can help accelerate transformation.


Enhance data capabilities. Enterprise data is still often fragmented, but it [needs to be prepared for AI](#) models so that in future,

agents can access up-to-date data in real time to allow continuous pricing of risk and continually updated insights.

Prioritize trust. Trust is the essence of the insurance industry and as a result, AI is likely to be rolled out across organizations more slowly than in some other sectors, so that the right governance guardrails can be put in place first. They will ensure AI operates safely, compliantly, and within budget through continuous monitoring.

And as AI agents become more sophisticated, insurers must also ensure transparency and accountability in the way agents work with external sources of data and outside organizations. Once the right governance is in place, the pace of adoption will accelerate in the years ahead.

The arrival of AI at scale across all sectors also brings pressure from customers to offer the interactions and rapid transactions they get from other businesses in more advanced sectors.





How to capture data and knowledge for humans and agents

- The insurance industry is transforming rapidly, creating a risk that critical data and knowledge will be left behind.
- Insurers must capture, consolidate, and operationalize data to stay competitive in a market being reshaped by digitalization.
- Developing agentic AI use in insurance is a cultural and training challenge: companies need to develop operating procedures, implementation procedures, and responsible guardrails to make the most of it.

Technology change and artificial intelligence (AI) adoption go hand in hand, and at the root of this joint challenge for insurers is an open question over whether the data that they use is fit for purpose.

Insurers have typically operated fragmented, unstructured, and incomplete data systems, with data siloed across locations and formats and knowledge and processes often undocumented. Against that backdrop, companies must now ready themselves

for deeper collaboration between humans and AI agents — AI tools that can operate autonomously to achieve specific goals — and compete with insurtech challenger companies that do not carry the burden of legacy technology. This means developing [strategic AI-ready data roadmaps](#) and capturing not just current data from proprietary and public sources, but also legacy data, institutional knowledge, and the how-to of procedures and processes to make themselves future-ready (Figure 1).

Figure 1. The Infosys AI readiness framework



Source: Infosys Knowledge Institute

Legacy tech, data challenges

The data challenge extends into mergers and acquisitions (M&A): although the pace is slowing, industry consolidation has left many large insurers grappling with [incompatible core systems](#). Historically, the primary aim of insurance M&A deals has been building size and influence, not technology advancement. The result is an industry marked by technological discontinuities, necessitating manual workarounds and dealing with the associated inefficiencies.

AI has altered how insurers deal with those issues. Crucially, it has also injected urgency into transformation efforts because it demands a [unified systems-and-data ecosystem](#) to function efficiently and comprehensively. Achieving this will require vision, compromises, and trade-offs.

If the technology opportunity alone were not enough, insurers are already seeing

intensifying competition from new digital-first financial businesses in areas including [embedded insurance](#). These new entrants have been growing slowly but steadily: Five companies in the Forbes Fintech 50 list for 2026 [were insurtechs](#), including cyber insurance provider Coalition. Keeping pace with digital-native businesses will only be possible with modernization.

Knowledge management

Companies benefit from effective knowledge management, but creating, sharing, and retaining information within a business is not always straightforward. Where data structures have been neglected or are inadequate, institutional knowledge often only resides in senior employees. In an industry that can skew toward older workers — research suggests 50% of the current UK insurance workforce [could retire in the next 15 years](#), for example — this rapidly becomes a problem. Indeed, statistics suggest a looming crisis in the US, where the rate of new hires will be unable to [keep pace with retirements](#). As experienced hires leave, knowledge of procedures and processes can be lost, undermining the business.

Studies of knowledge management involving insurance workers show they are often carriers of [knowledge](#) that is not recognized or codified as a business resource, but which has significant value in a rapidly changing corporate environment. This is distinct from the knowledge that is formalized and often recorded within documents. This brings obvious challenges both in terms of quantifying the scale and depth of this

knowledge, and then in gathering and [importing it into a data-first world](#).

Some insurers have adopted deliberate strategies to capture knowledge at risk due to staff turnover or retirement. These may be on a case-by-case basis. For example, when the vice president of underwriting at Canada's Special Risk Insurance Managers recently realized that one of her career mentors was [about to retire from a specialist role at the insurer](#), a staff member was appointed to document all that person's insights during the training and handover period in an attempt to record 20 years' worth of specialist experience that would otherwise be lost. PetSure in Australia, meanwhile, introduced a new and improved knowledge management system that made information [easier to find for everyone](#), reducing reliance — and pressure — on individual subject-matter experts. And Italy's Assinord Verona S.R.L. used employee interviews to understand more deeply how customers make [insurance policy-buying decisions](#) as part of its digital transformation strategy.

Right data, right domain

But there is no one-size-fits-all. Insurance businesses typically have multiple operational domains, and their data needs and AI potential vary. Product design, underwriting, pricing, claim, and policy administration are all markedly different in terms of underlying assumptions and procedures and bring different data-capture challenges. Nevertheless, best practice can already be found in established AI implementations in several of these domains.

For example, automation has often lagged in underwriting, where building or adapting insurance products has relied on unstructured data that has historically demanded manual review. However, Munich Re has developed a [generative AI underwriting tool](#) to develop insurance products in settings where legacy systems are too inflexible to respond cost-effectively. AI can also help detect fraud. Insurers have often used crude statistical approaches, which might be only partially accurate. To address this gap, Zurich Insurance is deploying a generative AI-based fraud-detection tool that supports investigations by analyzing available data, leaving its investigators more time to do what they do best — investigate. The tool scans for words often used in fraud claims, as well as common fraudulent indicators such as where several claims allegedly for different people are all filed from the same address. Live use is already increasing the number of confirmed fraud cases.

Meanwhile in claims processing, AI has often been framed as a way to combat rising costs. Although that need exists, the real opportunity offered by AI goes beyond this. For example, Aviva responded to escalating costs by integrating an AI tool that treated data as a strategic asset into [a wider transformation of the business](#). Aviva embedded AI across six business-wide dimensions (strategy, talent, agile operating model, technology, data, and adoption and scaling) so that it could build and implement AI tools across the entire claims function. So far, the company has cut the average time to assess liability for complex cases by 23 days and improved routing accuracy by 30%.


From AI tools to AI agents

As AI evolves at speed, companies are moving from thinking simply in terms of ring-fenced, limited-purpose AI to deploying AI agents. These are digital entities that use the multimodal capabilities of AI to accomplish and even set complex multistage tasks with some element of autonomy (Figure 2).

For insurers, this is particularly useful, as insurance processes typically involve a

lengthy chain of checks, confirmations, and decisions. AI agents can bring efficiency here. But insurers must also remember that agentic AI systems are constantly evolving, improving over time. This requires people implementing them also to be continually responsive to the potential, as well as the pitfalls, of the technology. Implementing agentic AI demands that organizations think about AI at a strategic level, asking questions such as: “What kind of business do we want to be?”; “What kind of workforce do we need?”;

Figure 2. Key components of agentic AI ecosystem



Source: Infosys Knowledge Institute

and “What kind of culture is best placed to be both receptive to AI and responsible in its use?”

One difficulty of designing an implementation path is that insurance businesses need to avoid the temptation of initiating sweeping transformations before having codified what works in test environments: a disciplined learning process is needed. Companies achieve success with AI more frequently after [transforming operating models and training the workforce](#) for AI, Infosys research has found.

AI itself can be used in this process. For example, First Mid Insurance, a large full-service US bank-owned insurance agency, used AI to transform its onboarding and training functions. Multiple acquisitions, often with distinct workflows and processes, had resulted in inefficient workflows, increased compliance risks arising from manual processes, and a 200-page training manual that was difficult for new employees to navigate. A knowledge retrieval system was created and all workflows automated. The

result? Faster onboarding, fewer errors, and [a 25% increase in team productivity](#).

Allianz, meanwhile, started smaller. It implemented agentic AI to automate and speed up processing of simple weather-related claims [in Australia](#). This allowed it to test the technology’s applicability and scalability across the wider business. The approach broke down agentic actions into discrete steps with human review of final outcomes, creating seven independent AI agents rather than delegating full process to a single software entity. Claims processing time fell from days to hours, and scalability was enabled through a [modular architecture](#) that will allow agentic AI to be rolled out for other relatively simple, high-frequency use cases including travel delays, simple vehicle claims, or property damage assessments.

The canonical data model

At the heart of every company’s potential to accelerate AI lies its data strategy. For many insurance companies, the desired end state of data modernization involves the creation

Old code routines may produce certain results, but the reasoning behind them may have been lost. Present-day users and developers might know what the code is doing, but not why.

of a canonical data model (CDM), a single source of truth that will obviate the costs and inefficiencies of a fragmented data legacy. The CDM concept itself is far from new — the model was first applied at scale in the 1990s — but it is becoming increasingly important to provide a solid, consistent information baseline critical for [successful AI deployment](#).

A CDM is a structure for data, not an AI entity or agent. The CDM is not a replacement for legacy databases, but rather a system for translating data from multiple legacy sources into the canonical form and a common language that an AI agent built on a large language model (LLM) can readily access. The lack of a canonical data strategy has been cited as the single most important reason why data transformations and [agentic AI implementations fail](#).

The language model may be an LLM such as those used in publicly available generative AI services like ChatGPT or Claude. In closely defined use cases, like those typical of insurance, an AI agent could be built on a small language model (SLM) trained on insurance-specific databases.

The advantages of LLMs are that they can

respond usefully to a wide range of set tasks and can handle a high level of complexity. By contrast, SLMs are domain-specific, can be developed in-house, and run on limited computing resources. Their reasoning is also more transparent.

A CDM can work with both LLMs and SLMs, but the essential requirement is that there should be a canonical data strategy, supported by a tested CDM process such as the [Infosys Policy Migration Solution](#). This uses [generative AI to aid the extraction and conversion](#) of legacy data to new file formats in new data repositories.

A final dimension of the data challenge takes us back to the issue of institutional knowledge. The content of legacy code does not necessarily reveal the intent. Old code routines may produce certain results, but the reasoning behind them may have been lost. Present-day users and developers might know what the code is doing, but not why.

This challenge is one that could be amenable to agentic AI-based solutions, as software developers are now experimenting with automated recovery of the [“lost intent” of legacy code](#). With agentic AI

Companies achieve success with AI more frequently after transforming operating models and training the workforce.

implementations, the migration of legacy code is managed by multiple agents with specific tasks, including orchestration of the migration workflow and transformation of source data to target data. But a third agentic element is added, where the AI agent correlates code against external documentation, such as regulatory sources and business process data. It seeks to determine why the code exists and how to realize that intent in the new architecture — or dispense with it, depending on human judgment. Experts suggest that overcoming this “code-to-context gap” can be critical to modernization, as well as retaining institutional wisdom.

Taking this approach, with human intervention hard-coded into the design of AI agents, could be critical to the future of AI in insurance, as well as in comparable businesses.

Today’s insurers must decide which path to take. Decades of M&A-led restructuring have not been matched by data and technology transformations, but AI advances demand new standards. The competitive risk of neglected transformation has become too great to ignore. Insurance companies need urgently to establish best practice in data migration and modernization, people-readiness, and agentic AI governance.

What happens next

To quickly and effectively capture data and knowledge across insurance processes, [stay](#)

[competitive, and preserve critical institutional knowledge](#), companies should:

First decide what they want AI to do.

It sounds simple, yet amid the pressure to select and implement AI quickly, many companies skip this important step in their transformation. Businesses should start by assessing the data intensity of potential use cases, how these fit into modernizing the legacy ecosystem, and how to ensure optimal AI collaborations between agents and experts across the business.

Prepare systems and people for a human-agentic future.

Foster understanding of the potential and limits of AI, and how these systems can interact most effectively, with governance placed at the center of training. At the same time, audit carefully where the organization’s crucial knowledge lives — in both humans and machines — to ensure institutional knowledge about processes and procedures is retained wherever possible.

Scale initiatives at an enterprise level.

But only after developing certainty over the business relevance of selected use cases, and the capacity of staff to work with AI agents effectively, efficiently, and compliantly.

Accept there is a new need for rapid

investment and transformation. There is always a cost to modernizing large complex data systems like those in insurance — but there may be significantly bigger costs, both practically and reputationally to the business, to not modernizing.



Ways to govern AI risk in insurance

- Regulators and insurers must work hard to keep pace with AI's potential to transform the industry
- AI tools are fast-charging productivity and profitability, but the shift from machine learning (ML) to large language models (LLMs) introduces new hazards such as probabilistic risks, where models could replicate inaccuracies, and new systemic threats like vendor concentration risk and adversarial attacks.
- Governance is emerging through diverse regional lenses, but insurers must remain compliant with existing regulations regardless of whether AI is used.
- As agentic AI begins to operate with a degree of independence, the lines between machine prediction and human judgment are blurring, making it difficult to assign legal responsibility.

The adoption of AI in insurance is evolving from narrow automation to agentic systems capable of underwriting, pricing, settling claims, and detecting fraud without human intervention. These systems can ingest and analyze vast amounts of data with a speed and granularity that humans cannot match.

Corporate leaders see the promise of a new era of efficiency, profitability, and client satisfaction.

Yet integrating LLMs and generative AI applications into existing actuarial science is not straightforward. The ability of LLMs

to extract and summarize information from unstructured data is a great help for drafting contracts and analyzing loss reports, for example. But generative AI tools also carry risks.

The upshot is an industry that is being urged to move fast, but safely. On the face of it, these are two irreconcilable goals. So, it is important to look closely at how insurers are integrating generative AI into their businesses, the guardrails they are putting in place for its safe development, and the regulatory approaches that are emerging to safeguard the industry's integrity. As insurers and their supervisors are discovering, the devil is in the details.

Fast build, uneven progress

Within the industry, the development and integration of AI tools have become synonymous with operational excellence. In the US, insurance group [Travelers](#) reports that more than half of all claims are now eligible for straight-through digital processing, leading to significant annual run rate cost savings.

Yet the picture is uneven. The [European Insurance and Occupational Pensions Authority](#) reports that nearly two-thirds of European insurers are already using generative AI, but most remain at a proof-of-concept stage. Infosys's own survey of [AI business value](#) in 2025 found that one in five AI use cases achieved no value or had been canceled.

Chubb's chairman [Evan Greenberg](#) puts it

plainly in a letter to shareholders: integrating AI is "iterative, gritty work. It takes time and sustained focus — hard yards." And he believes there are limits to automation. "Remember, when it comes to most insurance, people still want to deal with people," he writes. "It's a trust business."

New technologies, new risks

The [International Association of Insurance Supervisors](#) (IAIS) identifies three main concerns. First, the risk that AI models perpetuate biases in training data, leading to unfair outcomes in underwriting, pricing, or claims. Second, that increased reliance on AI heightens exposure to cyberattacks and data breaches. Third, the opacity of AI models, particularly those acquired from third-party vendors, makes it difficult to interpret decisions, validate outputs, and ensure compliance. At the heart of these issues lies the fundamental difference between conventional ML and LLMs.

Conventional ML is trained on structured data sets — proprietary policy-performance data enriched with external sources such as weather records, health statistics, and credit scores. Actuaries define the formulas, select the variables, and scrub the data of protected characteristics, such as age, sex, race, or income bracket, allowing full auditing of how a result is reached.

LLMs are different. Trained on unstructured data — text, images, video, and audio — they are a treasure trove of information for insurers, but they are likely to replicate inaccuracies or falsehoods in the original


training data. The [American Academy of Actuaries warns](#) that generative AI responses may be incomplete, out-of-date, hallucinated, or missing nuances required for actuarial work. In addition, LLMs could potentially fall foul of regulations because unstructured data can be used to bypass or infer prohibited categories of discrimination. For example, parsing someone's grocery shopping habits could be used to guess their ethnicity. And because LLMs are trained on constantly evolving sources, their results are not always reproducible, a significant headache for

compliance officers. [AI-related harmful incidents vary greatly](#), Infosys research has found, and most companies have experienced at least one (Figure 1).

Probabilistic versus predictive

Because LLMs are probabilistic systems, their outputs are expressed as likelihoods or a range of possibilities. They are programmed to predict the next likely token (a word or line of code) in a sequence. Actuarial models, by contrast, aim for precise, reproducible

Figure 1. Types of AI incidents experienced by enterprises



N = 1,502

Source: Infosys Knowledge Institute

outcomes — the same inputs yielding the same output every time. Yet there are clear advantages in integrating both approaches. Predictive models provide rigor; probabilistic models improve visibility, particularly over emerging risks.

This combined approach is valuable for the accurate pricing of climate risk, for example, where historical weather patterns are no longer a reliable guide to future losses. Instead of withdrawing from volatile markets entirely, insurers can use real-time data to dynamically underwrite complex risks that were previously deemed uninsurable.

Extreme granularity

Access to vast new data sets — geospatial data, sensors, aerial imagery, telematics, and behavioral data — is enabling individual- and micro-location-level underwriting for the first time. This granularity is transforming car insurance, where telematics, or the real-time measurement of driving habits, enables highly personalized pricing.

In an interview for this publication, [Vaibhav Anand](#), assistant professor at the Greenberg

School of Risk Management at St. John's University in New York, said the problem is that existing fairness and antidiscrimination rules "were not designed for this breadth and richness of data, nor for such fine-grained segmentation."

Granular AI underwriting is driving more accurate risk assessments and increasing insurers' profitability, but existing regulatory frameworks have yet to catch up with its potentially discriminatory impacts on policyholders. "Even if AI is fully compliant, it could end up pricing out entire micro-communities or systematically disadvantaging certain groups," he warns. Anand says the growth of granular underwriting calls for a new approach to regulation — one that includes screening for unfair outcomes as well as unfair inputs.

Blurred accountability

As the use of agentic AI becomes more widespread, regulators and the legal system are paying closer attention to [where responsibility lies when AI models](#) make mistakes. This is because the speed and compute power of AI systems is blurring



Antidiscrimination rules were not designed for this breadth and richness of the data ... Even if AI is fully compliant, it could end up pricing out entire micro-communities or systematically disadvantaging certain groups.

Vaibhav Anand

The Greenberg School of Risk Management at St. John's University

the lines between machine prediction and human judgment, making it difficult to assign accountability. Carissa Véliz, associate professor of philosophy and ethics at the Institute for Ethics in AI at Oxford University, has [argued](#) that AI recommendations are closer to commands than descriptions: the outputs of predictive AI may sound objective but are, in fact, normative in that they implicitly suggest decisions or courses of action. Professor Anand adds that the people coding these decisions at headquarters “may not be the doctors or underwriters who would normally exercise judgment in the field.”

Reliance on third parties

In May 2026, exams were disrupted at 9,000 universities and colleges across the UK, Australia, Canada, and the US, following a cyberattack on the educational software provider Canvas. The attack led Instructure, the maker of [Canvas](#), to “[reach an agreement](#)” with the criminals to return and delete the stolen data.


[Vendor concentration risk](#) is becoming a significant threat, and not just in education. Because building foundational LLMs from scratch is too costly, most insurers rely on APIs from a handful of tech giants. If a single dominant model experiences a prolonged outage, a structural security breach, or an unannounced weights update that shifts its reasoning logic, dozens of global insurers could see their automated underwriting and claims processing freeze simultaneously. This threat is now well in the sights of regulators.

Divergent regulatory response

A variety of approaches are emerging for AI governance in insurance. Regulatory fragmentation creates compliance uncertainty for global insurers and the risk of regulatory arbitrage, with companies locating to where regulation is lightest.

United States: rules under federal pressure

— The Trump administration’s December



2025 [National Policy Framework for Artificial Intelligence](#) aims at “global AI dominance through a minimally burdensome” approach, revoking an executive order from the previous president, Joe Biden, and establishing a task force to challenge prudential state AI laws. This puts federal policy at odds with the [National Association of Insurance Commissioners](#) (NAIC)’s [Model Bulletin on AI](#), adopted by [over half of US states](#) to enforce transparency and fairness standards. In response to agentic AI, some US states, including [California](#) and [Connecticut](#), have introduced bills demanding specific transparency and consumer rights protections against Automated Decision-Making Technology (ADMT). In Colorado, the [Senate Bill 26-189](#) would allow aggrieved individuals to request a human review of agentic AI decisions.

European Union: binding, risk-based legislation — The [EU AI Act](#) is the world’s first binding, horizontal AI law which applies to companies operating within the EU in all sectors. It categorizes systems on a scale from “minimal” to “unacceptable” risk. “High risk” systems, such as AI-based medical software or AI systems used for recruitment, [must comply with strict requirements](#) including risk-mitigation systems, high quality of data sets, clear user information, and human oversight. AI systems that represent a clear threat to fundamental human rights, such as “social scoring” by governments or companies, are considered an “unacceptable risk” and are banned.

For insurers, the most directly relevant high-risk categories are AI systems used for life and

health insurance pricing and creditworthiness assessment. These trigger requirements for human oversight, documentation, transparency, and conformity assessments. In practice, though, many AI use cases in insurance, particularly those involving automation of internal processes or lower-risk applications, are likely to fall outside the risk classification framework [and attract only minimal compliance obligations](#).

United Kingdom: principles-based and light-touch — Post-Brexit, the UK is encouraging a technology-agnostic, outcomes-focused approach through regulatory sandboxes. Both the Prudential Regulation Authority (PRA), which regulates around 1,500 financial institutions in the UK, including banks and insurance companies, and the Financial Conduct Authority (FCA) have [stated](#) they will continue to adopt a technology-agnostic, principles-based, and outcomes-focused approach. This makes the UK one of the most pro-innovation major jurisdictions for AI in insurance.

Asia: diverse and fast-moving — Singapore has developed a cross-sector framework based on fairness, ethics, accountability, and transparency ([FEAT](#)). This guides how algorithms should be designed and governed and how businesses should assess fairness, explainability, and accountability. Singapore’s Insurance Authority, for example, supports AI pilots and voluntary toolkits, like [AI Verify](#), that companies can use to test and demonstrate accountability. This positions Singapore as a hub for AI-driven insurance innovation with a relatively permissive but structured governance environment.

South Korea's [AI Basic Act](#), in force from January 2026, mirrors the EU's risk-based approach, while Japan is building sector-specific guidance through its Financial Services Agency, though currently with [no punitive sanctions](#).

Across the globe, however, the baseline is that insurers must comply with all existing regulations and laws governing the insurance sector, with or without AI tools. As NAIC makes clear: "Existing state insurance laws apply regardless of whether decisions are made by humans, algorithms, or third-party vendors."

Respond with education

Training initiatives are at the core of the insurance industry's response to emerging technological risks. As noted by Infosys CEO Salil Parekh, [workforce reinvention and reskilling are critical factors in AI success](#). Regulators require it, and the robustness of risk management systems depends on it. Under [Article 4](#) of the EU AI Act, AI literacy training for staff is a mandatory legal requirement.

In the US, NAIC expects insurers to maintain meaningful human oversight. These accountability structures only work if staff are trained to exercise that oversight.

As a result, insurers are treating AI upskilling as a strategic necessity. Programs like Singapore's SkillsFuture offer proven models and best-practice examples.

At the company level, major European

insurers, such as Allianz, AXA, and Zurich, have introduced comprehensive training programs, some in partnership with universities, covering the ethical and regulatory obligations of AI use, as well as its advantages. At [Allianz](#), AI training programs ranging from foundational AI literacy and generative AI ethics to topics like advanced prompting and ML had reached more than 144,000 employees in 70-plus countries by early 2026.

What happens next

Regulators and companies agree that effective governance and human supervision must scale with risk. High-risk applications (coverage decisions, claims denial, critical infrastructure) require rigorous bias testing, human oversight, documentation, and third-party audits.

Lower-risk applications, such as document processing, need a lighter touch. The [IAIS](#) says: "Insurers need to develop a combination of governance and risk management measures that are appropriate for their specific AI use case. For example, in certain circumstances the lack of explainability of a specific AI use case may be compensated by other measures such as increased human oversight and/or enhanced data management."

Yet despite the best efforts of legislators and supervisors, [regulatory frameworks do not yet capture the full range of risks](#) attending AI deployment.

As [Zurich](#) notes in a paper on safe AI transformation: "At the systemic level,

no comprehensive AI incident reporting framework exists, and concentration risk among a few large AI providers remains largely unaddressed.” To address these significant gaps, the following steps could help regulators and insurers create a safer environment for developing AI, while maintaining consumer trust.

Deploy an orchestration layer. To safely capitalize on agentic AI without breaking regulatory compliance, insurers must build robust infrastructure to govern autonomous agents. Designing and deploying an orchestration layer is one solution. This acts as an algorithmic air-traffic controller, explicitly defining data access parameters, monitoring agent interactions, and setting strict triggers for human intervention.

By demonstrating where an AI-assisted prediction ends and human judgment begins, it can help establish exactly who — and what — is responsible at every stage of a transaction.

Close the accountability gap. As agentic AI becomes more deeply embedded in insurance workflows, the law, legal theory, and regulatory frameworks will all need to

grapple with where responsibility lies when a chain of autonomous, non-human decisions produces a harmful outcome.

Gather better data on AI models. Better data on model reliability is essential. A standardized framework for reporting AI failures would not only build the evidence base underwriters need to assess AI risk, but also signal to regulators and consumers that the industry is serious about accountability.

Focus on outcomes rather than technology. To date, companies and regulators have concentrated on ensuring that the inputs of AI models — data quality, algorithmic objectivity — are in compliance. But given that technology will always evolve faster than regulation, a focus on outcomes is also required. Judging AI-enabled outcomes against well-established principles — fairness, freedom from bias, transparency — could be the basis of a durable, technology-agnostic regulatory framework.

Insurers who can demonstrate that their AI meets those standards will be better placed to innovate with confidence, and to maintain the trust on which the entire business of insurance ultimately rests.

Despite the best efforts of legislators and supervisors, regulatory frameworks do not yet capture the full range of risks attending AI deployment.

Buy, price, settle: The new arithmetic of AI ROI in insurance

- Insurance companies have committed to AI faster than they have learned to measure it effectively.
- A better approach maps AI investments to the defining outcomes of insurance economics: buy, price, and settle.
- Applying a value-focused framework can help put that principle into practice.

By February 2026, [nearly two-thirds of European insurers were actively using generative AI](#). Adoption is mainstream and investment is [already well-established](#) in agentic AI — systems that perform tasks autonomously with defined goals. That level of adoption suggests insurers find the technology broadly useful.


What most insurers cannot yet say is whether the investment behind AI is paying off. Infosys research shows that two in three insurers are in the testing phase or using AI only in isolated projects. Some 95% of respondents

working at insurers in the Infosys 2026 Future of Work survey said their companies are using AI. But 26% said that use was focused on the big task of transforming overall work and processes (Figure 1).

Meanwhile, increasing regulation, such as the EU AI Act, whose main regulatory provisions become enforceable on August 2, 2026, risks driving up costs and making the return on investments (ROI) calculation more challenging.

Still, insurance is a natural habitat for AI. In

Figure 1. AI use in insurance companies



N = 261

Source: Infosys Knowledge Institute, Future of work, June 2026

the experience of Infosys industry specialists, the majority of insurance work comprises document-heavy, back-office tasks. How does the industry turn a successful technology into one with a measurable return?

Right now, the board is asking whether AI is worth it and just gets back operational anecdotes. The fix is a framework that maps every AI investment to one of three outcomes that have always defined insurance economics:

1. How customers buy policies,
2. How the insurer prices risk, and
3. How it settles claims.

ROI math doesn't fit AI

Advocates of AI promise the technology will significantly increase revenue and reduce costs and risk, with organizations using it to rethink entire processes across their operations. Enterprises are using [AI agents to automate processes](#), speed up development, and improve customer service response. Yet most enterprises are still deploying it only to do existing work faster and more cheaply.

Only if they can measure gains convincingly will insurers get the most from AI. Yet the ROI frameworks insurers know best were built for a different kind of project. Cloud migrations and core systems modernization are bounded and substitutive; costs stabilize after deployment and value sits in one domain. AI breaks these assumptions. It sits on top of the legacy stack rather than replacing it. Its costs scale with use, and value spreads and compounds across functions.

British multinational insurer [Aviva spotlights](#) "over £90 million [\$120 million] of claims cost savings" from the AI models in its claims operations. No single model is responsible for all those savings, so no single measurement captures them. Aviva attributes the figure to a portfolio that spans claims summarization, fraud detection, and repair-network optimization; reaching a single figure from that spread is what forced the carrier to look at the technology impact as a whole.

Scope is one question; the type of value is another. Productivity and speed gains, for example, are often conflated. Speed is time spent on a task, while productivity is output

per unit of input. A 30% speed gain is not equal to a 30% productivity gain, because of bottlenecks between input and output: a task completed faster might just spend longer in a queue for human sign-off.

Boards have also started to tighten their timeframes. The proportion of insurance CEOs now expecting AI investments to pay back within three years has [grown to two-thirds](#), up from 21% a year earlier.

But tighter timeframes will not fix the problem. The difficulty of moving from pilot to production is a multifaceted problem, with one key challenge [being how value is measured](#).

Without a better framework, the industry faces a [technology overshoot](#), with the capability sprinting ahead of the business's ability to make use of it.

The cost of AI

While boards scrutinize return, cost is often underexamined. In the cloud era, companies shifted IT spending from one-off capital expenditures to ongoing but fixed, operational expenditure subscriptions. AI spending does not follow this model, because costs increase with use and AI vendors are shifting their pricing models.

The LLMs that underpin generative AI and many agentic tools are expensive to train and run, and vendors are growing impatient to recover their costs. In some cases, AI vendors are shifting from a per-seat license with unlimited use to a per-token or per-agent

model, where costs rise with consumption.

The risk is recurring bill shock: the spending behind the current AI boom will not be a one-off spend or a flat subscription like the last generation of technology. Instead, it will increase, so boards could end up under pressure to scrap projects that are delivering value.

Regulation adds further recurring costs. One [impact assessment](#) puts EU AI Act compliance at €29,277 (\$34,000) per AI system per year (€10,733 for robustness and accuracy, €7,764 for human oversight, €4,390 for documentation, €3,627 for information provision, and €2,763 for training-data compliance). That's without considering the risk of fines.

A further cost to the business is shadow AI. When the business moves too slowly to provide AI tools or offers tools that are less capable than the ones employees use on personal devices, some staff will buy their own tools. Aside from the cost of these extra subscriptions — and the security threat they pose — these tools will also increase IT fragmentation that will cost money to untangle later.

Buy, price, settle

Shifting to a framework built on concepts the board already understands will make AI ROI easier to measure. Within that framework, we can apply four metrics the industry already uses: productivity, accuracy, speed, and experience. Each subsection draws on these metrics, though the weights differ. Speed and

experience dominate buy and settle; accuracy dominates price.

Buy: How customers acquire policies

The buy-side metric emphasizes speed, productivity, and experience. AI that reduces time to quote is increasingly able to deliver value for the customer. Improving the ratio of total active insurance premium volume (in-force premium, or IFP) to full-time employees — a common performance metric in insurance — is a productivity boost. Finally, a higher net promoter score (NPS) after onboarding or a digital quote shows improvement in customer experience.

Lemonade's first quarter [2026 shareholder letter](#) supports claims of a productivity surge. The carrier ended the quarter with just over US \$1 million of IFP per employee, "having more than doubled IFP while reducing team size by 6% since the fourth quarter of 2022." IFP grew 32% year-over-year to US \$1.33 billion. (US insurers averaged year-over-year premium growth of 6% to 10% in the past three years, but most are starting with a much larger premium base in the hundreds of billions.) While Lemonade's

Cloud migrations and core systems modernization are bounded and substitutive; costs stabilize after deployment and value sits in one domain. AI breaks these assumptions.

technology-first approach allows it to operate its insurance business more efficiently, it has yet to achieve profitability. The company's customer acquisition costs, expansion efforts, comparatively small premium base, and high reinsurance costs have translated into a quarterly net loss since its July 2020 IPO.

Price: How to price risk

Better accuracy is the primary way to better price risk: sharper underwriting, more precise segmentation, and more exact actuarial models. One 2025 [analysis of AI insurance](#) estimates efficiency gains of up to 36% in complex commercial property and casualty (P&C) lines, up to 3% points of loss-ratio improvement from previously inaccessible unstructured data, and customer-service productivity increases above 30%.

An example of accuracy gains comes from Canadian insurer Intact Financial, [which reported](#) almost 600 AI models generating CAD \$200 million (\$142 million) in recurring annual benefits. The company expects that [figure to reach CAD \\$500 million by 2030](#). Few insurers quantify AI impact this precisely, but the example shows that AI ROI can be measured when the carrier emphasizes it.

Settle: How to pay claims

The final set of improvements comes on the settle side, affecting how claims are paid. This spans productivity, accuracy (indemnity, fraud capture, leakage), speed (time to first contact), and experience (NPS, complaint volume, hold time).

Allianz UK is an example of accuracy gains. It reported £37.7 million (US \$50 million) in fraud savings in the first half of 2024 through its [Incognito machine learning detection tool](#). Productivity gains are equally evident at Munich Re's ERGO unit, which expects AI to [replace up to 1,000 telephony and claims-processing positions](#) over five years. Though there will be no forced redundancies, a reskilling academy opening in 2026 will have capacity for up to 500 employees over two years.

Every AI investment in an insurer should be proposed based on one of these three outcomes. If a project cannot say which of buy, price, or settle it improves, it has no baseline and no metric, so it can never be proved to have paid for itself. That is the "unclear business value" Gartner warned about, and a potential cause of failure for AI projects.

The Infosys lens

The Infosys Value Realization Framework has been created specifically for insurance carriers. It is based on three factors: feasibility, desirability, and viability.

Feasibility asks whether the use case can be built within the carrier's current data, technology, and regulation. Desirability is about whether the target audience of customers, agents or underwriters would use it. Viability questions whether the line of business it touches can still earn money once the cost of AI is factored in. The factors run sequentially: if a use case is not feasible, then the other factors are not relevant.

Instead, insurers should start by picking the lever they most want to pull: growth, cost, or risk. Then they can develop three to five use cases per lever. These can each be deployed and the success of each tracked against a hard pre-AI baseline. They should be scaled up only if the baseline improves.

This framework guides people in choosing what to work on and tracking its return. Of course, there is a separate question: can the organization deliver the promised returns? This issue can be answered by a low-cost dry run before the budget is committed: the key stakeholders walk through a planned deployment, step-by-step, to identify where it might stall. One insurer ran that exercise and uncovered potential governance, compliance, and procurement issues, enabling it to plan the actual implementation better.

This sort of tabletop testing surfaces the gaps; closing the gaps needs an owner. That's the job of the chief AI officer (CAIO) — a role that [insurers, including AXA, have begun appointing](#) in recent years. Across all sectors, a study found that [76% of organizations have a CAIO](#), up from 26% in 2025.


It's easy to see why. Just [24% of insurance executives are very confident](#) they could pass an independent review of AI governance and controls within 90 days, even though 61% say their boards have set AI governance policies. The CAIO role exists to close the gap between policy and assurance, and the evidence suggests it is bearing fruit. In the survey above, companies with a CAIO spanned drinks, sports, and health, as well as technology (Figure 2).

Regulators are pushing the same way: the NAIC [recommends board-level accountability](#) on AI programs, and the UK's Financial Conduct Authority expects oversight allocated to a senior management function under the [Senior Managers and Certification Regime](#).

The responsible executive must leverage AI, but still needs people who can do the work. This requires a dual approach: augment the whole workforce with AI, and at the same time build deep AI engineering and domain expertise.

The most effective model is that of the [forward deployed engineer](#), someone who is both a domain expert and technologist. While they're not actually engineers, Aviva's 500-plus claims handlers, reskilled to wield the generative AI summarization tool, are analogous to forward deployed engineers. ERGO's 1,000-position restructure paired with

Figure 2. The rise of the CAIO role



Source: Data from IBM Institute for Business Value

its reskilling academy is a second approach. Full AI ROI will only be achieved if staff can be reskilled.

All of it runs to a single execution rhythm. The Infosys sequence for bridging pilot and production is Set Up, Scale, Evolve. The aim is to prove value in a contained first wave, scale where the baseline shifts, and evolve the estate as adoption deepens.

Vertical agents

Once insurers have an effective framework in place for measuring AI ROI, we see ways they can build value in the future. The first is by adopting “vertical agents” — autonomous AI tools built for the workflows of a single industry. In insurance, that means agents that handle underwriting, claims, or policy servicing, with the industry’s data and rules built in.

The momentum exists. Analysts suggest that AI and automation could improve expense ratios at the top 50 US insurers [by two points in 2026](#). The shift to agentic tools is happening more slowly. It is predicted that [fewer than 15% of organizations](#) will

turn on agentic features in their automation platforms in 2026. Insurers have stopped asking whether AI is useful. The question now is whether they can prove it pays, project by project, and act on what the numbers show. That is what will separate the insurers that lead from the ones left explaining their fruitless investments.

What happens next

Whether the industry leads or lags turns on whether boards can measure what their AI investments are returning. To get ahead, insurers should:

Evaluate every AI investment against buy, price, or settle metrics, and reject any that does not fit in a specific category. [Aviva’s AI claims tools](#) point to one outcome, enabling the carrier to credit more than 80 models with over £90 million in savings.

Choose the strategic lever targeted (growth, cost, or risk) before the use case. As the Infosys Value Realization Framework recommends, a named lever makes measurement possible and prevents projects sprawling.

If a project cannot say which of buy, price, or settle it improves, it has no baseline and no metric, so it can never be proved to have paid for itself.

Baseline pre-AI performance and track one primary metric per lever. Intact Financial shows that [rigorous measurement can identify trackable past value](#) and make reliable projections for future benefits.

Build a value bridge of leading indicators for work with a long payback. With [boards expecting quicker payoffs](#), AI deployments must clearly show likely value to survive budget cycles.

Appoint a CAIO or equivalent accountable owner. Companies such as [AXA are appointing CAIOs](#) to help ensure that projects meet governance requirements, as well as delivering measurable value.

Treat AI procurement as a board-level business line item with a named owner and forecast. AI costs recur and climb with use, so companies cannot just buy and forget.

Pressure test high-impact use cases with a structured dry run before committing budget. Governance or compliance problems are a common cause of failed AI projects; a dry run can help expose these before they become a problem.



Editors

Deepak Subbramaniam Associate vice president, Raleigh

Chad Watt Senior consultant, Dallas

Contributors

Infosys Knowledge Institute

Jeff Kavanaugh Vice president, Dallas

Infosys Insurance

Kannan Amaresh Senior vice president, Toronto

Karthik Balasubramanian Portfolio head, Basking Ridge

Atul Chaturvedi Associate vice president, London

Prashanth Dwarakanath Industry principal, Mysuru

Deb Dey Portfolio head, Hartford

Senthil K Vice president, Dallas

Roopa Navin Senior client partner, Boston

Dr. Hema Prem Raina Vice president, London

Anna Sanko Director, Raleigh

Jitin Sharma Group manager, Singapore

Vishal Sharma Group manager, Hartford

Infosys Consulting

Yogesh Bansal Partner, Newark

Madhu Thachappilly Senior principal, Chicago

Infosys marketing

Anshul Agarwal Marketing manager, Noida

Mark Iorizzo Associate vice president, London

Analysis and production

Sandeep Infosys Knowledge Institute, Bengaluru

Kate Bevan Infosys Knowledge Institute, London

Pramath Kant Infosys Knowledge Institute, Bengaluru

Isaac LaBauve Infosys Knowledge Institute, Dallas

Nikki Seifert Infosys Knowledge Institute, Dallas

Pragya Rai Infosys Knowledge Institute, Bengaluru

Pranav Tekade Infosys Knowledge Institute, Bengaluru

About Infosys Knowledge Institute

The Infosys Knowledge Institute helps industry leaders develop a deeper understanding of business and technology trends through compelling thought leadership. Our researchers and subject matter experts provide a fact base that aids decision making on critical business and technology issues.

To view our research, visit Infosys Knowledge Institute at infosys.com/IKI or email us at iki@infosys.com.



For more information, contact askus@infosys.com



© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.