

An original whitepaper by Keylane

# Innovation Built on Trust

Why the foundation and not the model,  
determines whether AI delivers value  
in insurance and pensions

---

With contributions from John Reynders, CEO · Nelson  
Petracek, CTO · Nicolas Marzin, Head of AI · Iris Reynisdottir,  
Head of Product



# Foreword

Every insurer and pension fund we speak with has invested in AI. Most have something running. But very few can point to tangible results, and boards have started asking **why**. It is a fair question.

The technology itself is not the problem. Today's models are extraordinarily capable, and they are improving so quickly that making a long-term bet on any one of them makes little sense. The problem is what it is built on.

The commitment is real. Boards are investing, and most organisations now have budgets and teams dedicated to AI. What decides the return is not the size of that investment. It is the foundation underneath it. Without trusted data and a direct connection to the systems where policies, claims and member records live, even well-funded AI struggles to move from promising to proven.

That explains why so many pilots remain pilots.

For AI to create real value in insurance and pensions, three things matter: the quality of the data underneath it, the governance around it and its connection to the systems that hold the record.

Get those right and the model itself becomes much less important. That is exactly what you want when models change this quickly.

Caution in our industry is understandable. I have spent my career in software, and I have never worked somewhere where the cost of being wrong lands so directly on people.

An AI error here is not just a software defect. It can affect somebody's claim, premium, or retirement.

But caution alone does not make AI safe.

Move too slowly and you risk carrying an outdated operating model into a market changing around you. Move carelessly and you risk automating poor decisions at a scale you cannot audit or explain.

This paper is about moving forward decisively, with AI that delivers real business impact. Safety is what makes that possible, not the goal in itself. We set out six principles we believe any AI vendor in this sector should be measured against, Keylane included. We explain why governance creates speed rather than slowing it down. And we show what we are building on the platforms that already run core administration across property & casualty and life & pensions.

Two decades of running those platforms has taught us that in this industry, what you deliver matters more than what you claim.

That is the standard we apply here, too.

**John Reynders**  
Chief Executive Officer, Keylane



# From adding AI tools to redesigning the work

The boardroom conversation about AI has changed. The question is no longer what AI can do. It is where the return is.

Most of what has been built so far was designed to prove the technology worked rather than create measurable business value. That is why we now see chatbots, summarisers and copilots sitting alongside processes that still work much as they did three years ago.

The bigger opportunity is to reimagine those business processes.

In property and casualty, claims, underwriting and servicing have long been candidates for event-driven automation. What changes is what happens after the trigger. A workflow no longer has to model every decision path in advance. It can reason about the specifics of a case against governed data, handle the variation that made these processes resist automation in the first place, and escalate to a person when the situation warrants it.

In life and pensions, AI can change how organisations deal with important member moments such as retirement, death, divorce, and transfers. Instead of waiting for the member to make contact, the organisation can respond when a relevant event occurs.

That is more than doing the same work faster. It changes what an insurer or pension fund can offer.



# The gap between activity and value

---

The research shows how wide the gap remains.

McKinsey reported in 2025 that 88% of organisations use AI in at least one function. Yet, only 39% can point to any effect on group operating profit, and for most of those, AI accounts for less than 5% of it.

MIT NANDA reported that 95% of enterprise generative AI pilots deliver no measurable impact.

Those figures do not necessarily describe technology that failed. They describe organisations that ran experiments without fundamentally changing how the work gets done.

The organisations seeing real impact are changing processes, not simply running more pilots.



# Starting at the wrong end

---

Many AI programmes struggle to move beyond experimentation because they start with the technology.



What I see is a technology-first approach. It is almost a race. Organisations buy tools and platforms and are still looking for problems to solve with them. Everybody needs to try the latest model. If you start from the technology and go looking for use cases, you start small, your ambition is low, and the value is small.

---

**NICOLAS MARZIN**, HEAD OF AI, KEYLANE



The better starting point is the problem worth solving. Then decide whether AI is the right technology for it.

---

"The value of an AI subscription on its own is zero. It is only worth something when you apply it to a specific business problem, and then the value is proportional to the ambition."

**NICOLAS MARZIN**, HEAD OF AI KEYLANE

Sometimes AI will be the answer. Sometimes a simpler process or better user experience will create more value.

Programmes also stall because people do not trust AI output, do not know how critically to assess it or have not been given the skills to use it well. A better model does not solve those problems.

Scaling too early adds another risk. One data leak or consequential decision without proper oversight and an audit trail can undermine confidence in the entire programme.

---

# What makes insurance and pensions different

---

Machine learning itself is not new to insurance. Dynamic pricing, fraud detection and next best action have been used for years.

Generative AI is different.

Its output is probabilistic. The same input does not always produce the same answer. Training data can be opaque. Bias and hallucination are documented characteristics of the technology.

And in this sector, the output can ultimately affect what a claimant is paid or what a member retires on.

So the critical question is not which model to choose. It is how to use a probabilistic technology in an environment where purpose, authority, and boundaries are clearly defined.

What may the AI do? Which data may it access? How is its output monitored? When does a person need to intervene? Can the organisation show afterwards how a decision was reached?

Those questions are answered through governance and orchestration, not model selection.



# A sector that understood the risk, then got stuck

---

EIOPA surveyed 347 insurance undertakings across 25 countries in 2025. It found that 65% were already using generative AI and a further 23% planned to within three years.

But most deployments were still at proof-of-concept stage. 64% of use cases focused on back-office productivity, while 83% retained human input or validation. Privacy, security, regulatory compliance and skills were among the barriers mentioned most often.

There is no equivalent survey for pension funds, but many of the same governance questions apply.

The caution makes sense. What is missing is a clear route from controlled experimentation to controlled deployment at scale.

Governance is what makes that route possible.



# Why governance accelerates rather than constrains

---

Governance, risk and compliance are often seen as the brakes on AI. That reputation comes from doing them last.

Teams experiment and deploy until someone asks about certification, regulation, security or the evidence a supervisor will expect. Then the programme stops and the rework begins.

Governance gets blamed for the delay. In reality, the delay came from addressing it too late. Speed without governance is a dangerous blind spot.

Build controls in from the start and there is much less to unpick later.

“



The choice is not between speed and safety. With the right architecture, trust becomes the source of speed. The riskiest strategy is uncontrolled experimentation. The second riskiest is waiting for certainty while the market moves around you.

---

JOHN REYNDERS, CEO, KEYLANE

There is a useful parallel with the development of the car. The technology alone did not make driving something billions of people would trust every day. Roads, traffic rules, speed limits, and safety measures developed around it.

“

Governance did not kill the car. It turned it into something billions of people trust their lives to every day. It is not a constraint on ambition. It is what makes ambition trustworthy enough to scale.

---

IRIS REYNISDOTTIR, HEAD OF PRODUCT, KEYLANE



# Reading the EU AI Act as design discipline

---

From August 2026, AI used for risk assessment and pricing in life and health insurance is classified as high-risk under the EU AI Act.

Treat the Act as a checklist and it creates work. Treat it as design discipline and many of its requirements are straightforward: understand the use case, assess the risk, apply appropriate controls, be transparent and keep evidence of what happened.

Insurers and pension funds already work within frameworks including Solvency II, IFRS 17, GDPR and DORA. Evidencing decisions is not new.

AI increases the number and speed of decisions that may need to be evidenced. That is why the answer needs to be structural.

If the platform creates that evidence while running the process, compliance becomes part of the system rather than a separate exercise afterwards.



# What good looks like

---

AI will soon make decisions in insurance and pensions at a scale no team can check by hand. Whether that becomes an advantage or a risk depends on one thing: trust.

Core administration platforms already earn that trust every day. When they stop, claims and payments stop too. AI working inside those processes has to meet the same standard.


That standard comes down to three things. Know it. Control it. Keep it safe.



“AI can act, but only within a defined mandate. Every agent has to be known, owned, classified, guard railed, observable, and reviewable. An AI error is not merely a software defect. It can affect a claim, a premium, or a pension outcome. Somebody’s financial security.”


---

JOHN REYNDERS, CEO, KEYLANE



## 6 PRINCIPLES

We believe there are six principles every leader should insist on before approving AI infrastructure for a regulated business.



## 01 Governed data

AI needs access to trusted data, not uncontrolled copies, extracts and spreadsheets.

The organisation needs to know which data a decision used, what it means, and who was allowed to access it.

Data quality and ownership were challenges before generative AI arrived. AI makes solving them more urgent, and also necessary.

## 02 Auditability

A regulated organisation needs to be able to reconstruct what happened.

Which model ran? Which data did it use? Which workflow triggered it? What output did it produce? What happened next? Can it be replicated?

If that evidence is scattered across multiple solutions, governance gets harder with every use case.



## 03 Controlled authority

An AI agent should never have undefined power. Its purpose, owner, permissions, boundaries, and escalation points need to be clear.

## 04 Meaningful human oversight

Putting a human in the loop does not automatically make a process safe.



As humans are asked to be in the loop, over time they take everything to be normal. Involving a human in the loop becomes a check mark rather than a governance step.

---

NELSON PETRACEK, CTO, KEYLANE

Ask someone to approve hundreds of AI outputs every day and approval can replace genuine review.

Oversight therefore needs to reflect the decision. Some steps can run automatically. Some need four-eyes control. Others need expert judgement.

A claims handler, for example, may decide to pay a claim that is not strictly required under the policy wording because the customer relationship matters more than the amount involved.

That decision should stay human.

The question is not whether a person is involved. It is where human judgement genuinely matters.

## 05 Model independence

Today's leading model may not be tomorrow's. Models improve. Prices change. New providers and vendors emerge.

An organisation should be able to change the model without rebuilding the process around it.

## 06 Runtime safety

AI needs to stay safe after deployment, not just at the point of approval. This requires continuous monitoring for harmful content, jailbreak attempts, hallucinations, PII exposure, and task drift.

An AI control tower should surface early warning signals, trigger alerts and make suspicious behavior visible before it becomes a bigger problem.

And when an agent behaves outside its mandate, there needs to be a clear way to stop it. A kill switch is a basic control for AI operating in regulated, mission-critical environments.



# What Keylane is building, and why

---

Keylane is building Keylane Automate: a governed AI layer on the core platforms insurers and pension funds already run on.

We have made this kind of move before. We committed early to cloud and SaaS, when many expected core administration to stay on-premise. Today, 37 million policies and members run on the Keylane cloud.

We are approaching AI the same way: dedicated development streams, investment in people and partners, and production and governance capabilities built rather than announced.

It also means working with customers on early deployments across property and casualty, life and pensions, starting with AI applied to document handling inside existing processes.

We are using AI for our own product development and delivery too. We cannot credibly ask customers to change how they work if we are not prepared to do the same ourselves.

## The strategy underneath it

Our view is straightforward: language models are already becoming a commodity, and the model will not be the most valuable layer in the AI stack.

Models are changing too quickly. Building a long-term strategy around one provider means tying that strategy to the fastest-moving part of the stack.

The more durable value sits elsewhere: trusted data, domain knowledge, permissions and business processes.



Our bet is that Keylane will be the company that makes AI useful in a safe, scalable way for insurance and pensions.

---

JOHN REYNDERS, CEO, KEYLANE



# Reimagine your business processes

A core platform holds the truth: policies, claims, products, members, rules, permissions, and definitions.

But a traditional system of record largely waits for someone or something to initiate an action. Workflows are often fixed and changing them can take time.

That is the problem Keylane Automate is designed to address.

THE SAME GOVERNED DATA



DOING A DIFFERENT JOB

It turns a system of record into a system of actions. Customers build workflows and agents on governed, auditable insurance and pension data, without writing code and without waiting for a release.

IRIS REYNISDOTTIR, HEAD OF PRODUCT, KEYLANE



The data does not move. The governance does not loosen.

What changes is how teams can act on that data.

They can design, test and deploy workflows and agents across claims, underwriting, member communications and case handling, including processes that span the core platform and other systems.

Automate sits above the system of record rather than requiring changes inside it. The core continues to do the job it was designed for.

# What the architecture does differently

---

Keylane Automate includes a workflow designer, a catalogue of approved AI agents and models, end-to-end lifecycle governance covering development, risk classification and cataloguing, an AI Control Tower for safety and compliance, and connectors into the core and customers' own systems. Nothing runs until it is approved. The difference lies in three architectural choices.

## The workflow is the only entry point

An AI asset can be used through a governed workflow. Before entering that workflow, it needs to complete the approval lifecycle.

Governance happens through the process itself rather than alongside it.

### AI reads governed core data

---

Keylane Automate sits on the platforms that hold the record.

Agents can work within the roles, permissions and definitions already in place and operate within European data residency.

That reduces reliance on copies and extracts, where lineage and access control can become harder to manage.

### Models are catalogued, not hard-wired

---

Agents and models are registered, approved and managed through the AI Hub.

Organisations can therefore change models without rebuilding the workflows around them.

Point solutions can work well for individual use cases. The challenge comes as the number grows. Each may introduce its own data flows, approvals and logging.

One governance model avoids rebuilding those controls every time.

# Three levels. Start at any of them

AI can play three different roles in a process. Organisations do not have to work through them in order.



## **Embedded AI:**

summarisation, extraction or drafting within screens employees already use. These capabilities are useful, relatively low-risk and require little change to the underlying process.

## **Task-level agents.**

AI performs a specific task within a familiar process, such as checking coverage or gathering information a handler would otherwise collect manually. This is where measurable operational value becomes easier to see.

## **Orchestrated, AI-native workflows.**

Agents work with other agents and people across an end-to-end process. Claims stop following a path modelled in advance and start adapting to the case in front of them. Member service acts on a member's circumstances rather than waiting for a request. Workflows can move between the core platform and the wider technology estate.

If each level needs its own tooling, integrations and governance, every move becomes another technology project. One platform for all three means you can start where the value is.

If a customer wants features embedded in the platform, Automate provides that. If they are moving beyond task-level agents into orchestrated, AI-native workflows, Automate is built to support that too. Customers have a journey with us, just as we had a journey with the technology.

NELSON PETRACEK, CTO, KEYLANE



# Where the value lands first

---

In property and casualty, the first business cases are likely to sit in high-volume, repeatable work.


Claims are an obvious example. Where a handler currently touches every claim, a governed workflow could allow straightforward claims to move through automatically while requesting missing information directly from the customer.

That creates opportunities for straight-through processing and leakage reduction while keeping people focused on cases that require judgement.

In life and pensions, one of the more interesting opportunities is proactive service.

Spotting the event was never the difficult part. Working out what it means for this member, with this product, at this point in their life, is what required a person. That assessment can now happen inside the workflow, with a person brought in where the decision warrants it.

That gives life insurers and pension funds the opportunity to serve people when something changes rather than waiting for them to make contact.



# The question worth asking

---

There is a question we use internally to challenge assumptions about existing processes.



What would happen if you had a completely AI-automated product? If I was going to offer an insurance product, what if it was actually zero touch, with no human involved, ever? Now, that is probably too much of a stretch. I do not know if we will ever get to that point. But you almost have to think about it that way to see where and how AI can be leveraged.

---

NELSON PETRACEK, CTO, KEYLANE

The point is not that insurance and pensions should become entirely automated. For many products and decisions, they should not.

The question forces you to identify where human involvement actually adds value.

Which steps need judgement? Which need accountability? Which still involve a person simply because the process was designed that way?

That is a different starting point from looking at today's process and asking where AI could save a few minutes.



# The decades problem

---

During that period, a member moves, changes jobs, marries, divorces, has children and retires. Their circumstances change continuously.



You are going to have the situation where AI along the way is going to have to augment and facilitate that person's journey with those products over potentially many years.


---

NELSON PETRACEK, CTO, KEYLANE

The opportunity is not simply to complete a transaction faster. It is to maintain continuity of attention across a relationship lasting decades, and across millions of members.

No human team can hold all of that context continuously at that scale.

Noticing the change was never the hard part. Interpreting what it means for this member, against their history, their product and where they are in a fifty-year relationship, is what took judgement and time. No team can do that for millions of people while it still matters. That is the gap AI closes here.



# What changes underneath

The technology will keep evolving.

We expect more small, domain-specific language models built for particular insurance and pension tasks, alongside knowledge graphs, semantic layers and new approaches to memory.

The model underneath a workflow will change too, probably more than once.

That is the practical argument against hard-wiring today's preferred model into a long-term architecture.

The interface will change as well. Generative AI spread quickly, partly because people could simply type what they wanted. That has reset expectations of how customers and employees interact with software and data.

The answer is unlikely to be a chat box added to every portal.

The bigger opportunity is to connect natural language text or even voice, safely to governed data and processes.

Roles are going to change within our customer base. There is likely going to be new roles that show up and other roles that have to transform. It changes the way in which you think about offering capabilities to your customers, just like we have had to.

**NELSON PETRACEK, CTO, KEYLANE**



Customers will also find uses we have not predicted.



What excites me is that some customers will come with unexpected ideas for workflows and use cases we have not thought about, and maybe drive the product in a different direction.

**NICOLAS MARZIN, HEAD OF AI, KEYLANE**

That is one reason the platform needs to be composable.

# The ambition

Our ambition is for Keylane to become the standard for trusted AI across European insurance and pensions.

That has to be earned.

The work now is practical: turning early projects into repeatable solutions, proving the return, delivering production-grade reliability and keeping the architecture open as the technology underneath it changes.



# Where to start

---

Do not start with technology and look for a problem.

Start with the business value.

Choose one strategically important workflow where both the customer outcome and the economics are meaningful. Understand how it performs today. Redesign it end to end. Decide where accountability sits and what needs to be auditable.

Then prove the value before scaling.

Trying to do everything at once creates a lot of AI activity and very little impact.



# An invitation

Keylane brings two things to this work: deep knowledge of insurance and pensions, and core platforms already running critical administration across the sector.

Choose the use cases that can make a material difference. Redesign the workflows. Put governance in place from day one. Prove what works, then scale it.



If we act from a foundation of partnership, both organisations learn and develop at a speed they have never experienced before.

JOHN REYNDERS, CEO, KEYLANE



A trusted foundation you can prove.  
The freedom to build on it.

Talk to Keylane about Keylane Automate:

[keylane.com](https://keylane.com)

# In their words

---



The winners will transform workflows, not add AI tools. An AI error is never just a software defect here. It is someone's claim, premium, or pension. So, we lead with control, and we intend to lead.

---

**JOHN REYNDERS**  
CHIEF EXECUTIVE OFFICER, KEYLANE

Keylane has run the core for insurers and pension funds for years: the data, the rules, the workflows. Automate is the natural next layer. Governed innovation is how we have always had to build.

---

**NELSON PETRACEK**  
CHIEF TECHNOLOGY OFFICER, KEYLANE



For years, customers had to choose between speed and safety. We are building the platform that gives them both. One governance model, every asset, design time, and runtime. Then we hand them the keys.

---

**NICOLAS MARZIN**  
HEAD OF AI, KEYLANE

Governance did not kill the car. It made it something billions of people trust every day. Build the seat belts in from the start, and AI in insurance and pensions can go as far as anyone dares to take it.

---

**IRIS REYNISDOTTIR**  
HEAD OF PRODUCT, KEYLANE





**McKinsey & Company**, *The State of AI, 2025*. 88% of organisations use AI in at least one business function; 39% report any EBIT impact, most below 5%.

**MIT NANDA**, *The GenAI Divide: State of AI in Business, 2025*. 95% of enterprise generative AI pilots deliver no measurable impact.

**EIOPA**, *Generative AI Market Survey, August 2025* (EIOPA-BoS-25-679). 347 insurance undertakings across 25 countries, covering approximately 80% of EU gross written premiums. 65% currently using generative AI, 23% planning within three years; 64% of use cases back-office; 83% retain human input or validation.

**European Union**, *Regulation (EU) 2024/1689 (the AI Act)*. High-risk classification for AI used in life and health insurance risk assessment and pricing applies from August 2026.

Quotations are drawn from interviews conducted in August 2026 and edited for length and clarity.

