

Why AI won't scale on legacy life insurance systems.

And how to make it work.

A framework for scaling
AI in life insurance operations

Executive Summary

Artificial intelligence is high on the agenda of most life insurers right now. Boards are asking about it. Exec teams are playing catch up. Pilots are running across underwriting, claims, and distribution. On paper, momentum looks pretty good, but in most organisations the operational core of the business is largely unchanged.

That disconnect matters more than it might appear. In life insurance, long-term results aren't just about acquisition and claims outcomes - they're shaped by the efficiency of in-force servicing, billing, payments, and customer operations. These are high-volume, rules-intensive domains where cost-to-serve accumulates quietly, year after year. They're also where most AI initiatives haven't gone yet.

A useful way to see the scale is to think in interactions, not technology. Over a typical policy lifecycle (around 22 years), each life policy generates roughly 14 manual servicing events (tasks related to renewals, payments, reporting activities, and alterations). That doesn't sound dramatic until you multiply it across a real book of business. Across one million policies, that's around 14 million operational interactions an insurer must process over time. Across several million policies, it becomes an operating reality that dominates how organisations are staffed, how systems are configured, and where complexity quietly compounds.

This is why many AI initiatives to date have delivered visible activity without shifting the underlying operating model. Most have focused on augmenting decision-making at the edges - document summarisation, communication drafting, call centre assistance, underwriting support. These aren't trivial, but they don't address the architectural and workflow constraints inside legacy core systems. More importantly, they don't materially reduce the volume of repetitive servicing interactions that consume operational capacity.

So, the strategic question isn't simply which AI use cases to pursue. It's whether the operational core is structurally ready to support AI at scale and become economically significant. This paper argues that meaningful AI transformation in life insurance requires moving beyond "edge AI" toward structural readiness at the operational core. Without that shift, insurers will find themselves investing heavily in AI capabilities that remain bottlenecked by architecture.

Insurers that pull ahead in the next decade will not be those that experimented with AI tools or launched the most pilots. They will be those that evolved their operational foundations so AI could move from augmentation to orchestration. Moving from assisting individual tasks to reshaping the operational workload itself - reducing the millions of small, repetitive servicing interactions that quietly define life insurance operations today.

The illusion of AI progress

There is no shortage of AI activity in life insurance. Boards are receiving updates on generative AI pilots. Technology teams are showcasing copilots, chat interfaces, and automation tools. On the surface, progress looks rapid. But in most insurers, the underlying operating model hasn't moved.

Policy servicing workflows are still distributed across multiple systems. Billing and payment processes remain batch-oriented and reconciliation-heavy. Operational data is fragmented across legacy cores, ancillary platforms, and spreadsheets. Business rules are embedded in code, documentation, or the heads of people who've been there for twenty years.

In this environment, AI gets applied as an augmentation layer - summarising documents, drafting communications, helping call centre staff, accelerating analysis in underwriting and claims.

These initiatives can deliver measurable efficiency gains. But they don't alter the structural drivers of cost-to-serve.

In life insurance, cost accumulation is gradual and systemic.

It comes from:

- High volumes of policy alterations and endorsements
- Payment retries and arrears management
- Compliance-driven communication complexity
- Reconciliation between siloed systems
- Manual exception handling across edge cases
- Closed-book technical constraints

These dynamics are embedded in workflow design and system architecture. AI applied at the surface can reduce friction - it doesn't remove structural inefficiency. There's a real risk, then, of confusing visible AI adoption with genuine operating model transformation.

An insurer can deploy multiple AI tools and still find that servicing headcount remains high, payment complexity persists, exception handling continues to dominate operational effort, and unit economics are largely unchanged. That's not a failure of AI capability, but a reflection of where AI is being applied.

To materially shift cost-to-serve in life insurance, AI needs to interact directly with the operational core - where policy data, payment logic, business rules, and workflow orchestration converge. Without that, AI initiatives scale in number but not in economic impact.

Where cost-to-serve actually lives

Life insurance brands are built in moments of service, not moments of sale. The clarity of a communication during a premium increase. The consistency of a decision during a policy alteration. The confidence conveyed during a claim. The absence of friction when a payment fails.

These moments are operational. They're shaped by workflows, data integrity, and system design. For that reason, cost-to-serve in life insurance isn't purely a margin metric - it's directly linked to trust.

In practice, cost-to-serve is rarely driven by a single inefficient process. It's the cumulative effect of structural complexity embedded across the in-force book. Three structural pressure points drive disproportionate cost.

Closed books and product generational drift

Most established life insurers are running multiple product generations simultaneously. Products get repriced. Benefits are redesigned. Regulatory requirements evolve. Distribution strategies shift. Policy terms are refined. Older generations close to new business but stay active for years - often decades.

The result is what I'd call product generational drift: different rule sets, different premium structures, different contractual definitions, different data schemas, often sitting on different system representations. Servicing teams aren't managing a coherent product environment. They're managing a layered historical archive of design decisions.

Closed books introduce:

- Parallel workflows
- Embedded legacy logic
- Manual overrides to accommodate historical conditions
- Data structures that no longer reflect current product thinking

When those product generations sit on separate cores or partially integrated systems, the complexity compounds. AI can summarise policy documents or draft communications across these books. But unless the underlying product logic is unified (or at least coherently exposed) AI inherits the fragmentation rather than resolving it. Closed-book management isn't just a technology inconvenience; it's a root cause of cost-to-serve.

Payments as an operational multiplier

Payments and billing look mechanical from the outside. They're not. They're tightly coupled to lapse management, arrears handling, reinstatements, customer communication cycles, and compliance reporting. In many environments, policy and payment systems aren't deeply integrated. Payment failures cascade into servicing queues. Arrears logic gets layered over time. Reconciliation processes become permanent fixtures.

The cost isn't just transactional; it's behavioural. Payment friction increases lapse risk and generates servicing demand that wouldn't otherwise exist. Where payments are architecturally separated from policy logic, AI can help individual operators. It can't easily re-orchestrate the workflow end-to-end.

Exception density in servicing workflows

Life insurance doesn't run on the 'happy path'. Servicing teams deal with non-standard alterations, incomplete documentation, edge-case benefit interpretations, transitional states between policy conditions, and historical decisions requiring contextual judgement. Automation handles predictable flows well. High exception density, though, usually reflects structural product and system complexity, not just process inefficiency.

AI assistance can reduce handling time per case. It doesn't automatically reduce the number of structurally induced exceptions. That requires something deeper.



Three levels of AI maturity in life insurance

AI adoption tends to get discussed in binary terms - either an insurer is “using AI” or it isn’t. The reality is more graduated. In practice, there are three distinct levels, and the gap between them is significant.

Level 1 - Edge AI

AI is applied to discrete tasks at the periphery of core workflows: document summarisation, call transcription, internal assistants, underwriting decision support, claims triage, communication drafting. These use cases deliver productivity improvements but operate entirely within existing workflow boundaries.

Most insurers globally are here. Some are deeply invested at this level, which can create the impression of AI maturity without the underlying architecture to support it.

Level 2 - Workflow AI

AI becomes embedded within defined operational domains. It begins to orchestrate multi-step processes, trigger downstream actions, integrate structured and unstructured data, support exception routing, and provide contextual decision support tied to actual policy state.

This requires deeper integration with core systems and more coherent operational data. It represents a meaningful step up from Level 1, but still operates largely within the constraints of existing system architecture.

Level 3 - Structural AI

At this stage, AI becomes integrated into the workflows that manage the in-force book. When operational systems expose consistent data and reliable workflow triggers, AI can operate within the flow of work rather than alongside it.

It can guide operators through complex cases, identify patterns such as lapse risk, route exceptions, and generate communications based on policy context. At this level, AI begins to influence how operational work is structured, not just how individual tasks are performed.



Most insurers believe they’re operating at Level 2. In practice, the majority are still at Level 1. The integration depth required for genuine workflow AI is more significant than it appears from the outside.

The progression to Structural AI is not primarily a model challenge. It’s an architectural and operating model challenge, and that’s exactly what makes it hard to prioritise when there are faster wins available at Level 1.

The architectural constraint

AI systems perform best in environments with clear domain boundaries, coherent data models, explicit business rules, real-time processing, modular integration surfaces, and workflow orchestration capability. Legacy core platforms were built for durability and compliance. Composability wasn't the design brief.

Over time, additional layers accumulate: integration middleware, servicing overlays, data warehouses, reporting tools, automation scripts, manual workarounds. Each one individually rational. Collectively, they produce an environment that AI finds genuinely difficult to operate in. Fragmented data, embedded rule logic, batch-driven workflows - these aren't just technical inconveniences. They're the reason AI stays at Level 1.

Structural AI requires deliberate, sustained evolution of the operational core

That means:

- Rationalising product generations
- Exposing business rules through structured interfaces
- Consolidating data models
- Introducing event-driven processing
- Reducing reconciliation dependencies across systems

AI capability scales in proportion to architectural integrity. That's not a principle or a theory - it's what you observe when you're working inside these environments.

The risk of building tomorrow's legacy

AI investment is accelerating. New tools are introduced. Vendors extend capabilities. Internal teams run pilots. Individually, most of these initiatives can deliver value. But if they're not structurally aligned, they can collectively increase architectural surface area rather than reduce it.

AI layered onto fragmented environments tends to introduce additional integration dependencies, an expanded vendor ecosystem, parallel logic paths, and increased governance overhead. Over time, insurers risk creating an "AI overlay" architecture - intelligent tools coexisting with legacy workflows rather than reshaping them. The tools get smarter. The operating model stays complex.

The alternative path - pursuing core evolution with structural intent - looks slower in the short term but produces compounding returns. Insurers that take it can start to rationalise workflows across product generations, expose reusable business logic, consolidate data representations, and design AI capabilities that operate consistently across domains rather than solving the same problem repeatedly in different contexts. The long-term advantage lies not in any single AI model, but in the coherence of the system that model operates within.

From system of record to system of action

Life insurance exists for one purpose: to pay claims and support people at moments of vulnerability. Everything else - the technology, the processes, the operational design - exists to make that promise sustainable over time.

When AI is embedded thoughtfully within the operational core, it can reduce friction in servicing, detect avoidable lapses before they happen, surface relevant context in complex cases, improve the clarity of communications, and free professionals from the kind of repetitive work that adds no value to the customer or the business.

The goal is not to remove human involvement. It is to reduce fragmentation so human effort is applied where it matters.



A system of record stores information. A system of action orchestrates it. The evolution from one to the other requires deliberate core evolution - reducing fragmentation, clarifying boundaries, aligning data and rules with how insurers actually operate.

The true measure of AI maturity in life insurance won't be the number of pilots launched or the size of the AI budget. It will be whether, five years from now, the people working in these organisations find their jobs less fragmented, operational processes run with far greater efficiency, and customers experience service that feels simpler and more responsive in the moments that matter.

About the Author




Tim von Dadelszen
Chief Product Officer at Simfuni

Tim works with life insurers on modernising operational platforms supporting in-force policy management and payments. His perspective is shaped by the practical realities of core evolution within mature life insurance books - what works, what stalls, and why the gap between AI ambition and architectural readiness is the defining challenge of the decade ahead.

About Simfuni

Simfuni is an AI-powered platform for life insurance operations, orchestrating in-force servicing and payments by unifying core policy and billing systems with customer interactions and workflows.





Orchestrating modern insurance operations

✉ For more information or permission to reprint, please contact info@simfuni.com
in To keep up to date with our latest news and content, please follow us on [LinkedIn](#)

simfuni.com

© 2026 Simfuni. All rights reserved.