

SHIFT



Beyond the pilot: the IT leader's guide to production-grade agentic AI

www.shift-technology.com/en-gb

Introduction: the build vs. buy paradox

Every insurer with a capable engineering team is currently weighing up a critical decision: why pay for a specialised platform when we can build on top of foundational models ourselves? The allure of DIY is strong; general-purpose large language models (LLMs) are genuinely impressive in their ability to reason, summarise, and generate. However, the transition from a successful demo to a scaled, auditable production environment is where most in-house initiatives stall.

For IT leadership, the challenge isn't just about technical capability—it's about navigating the hidden costs of regulatory weight, legacy integration, and the "PoC Purgatory" that traps the majority of internal projects (Holland, CIOdive, 2026). Today, fewer than half of insurance businesses have deployed AI in even a single function, with end-to-end workflow automation being the rarest of all. This paper explores the critical concerns dominating the 2026 IT agenda and why a purpose-built architecture is the only viable path to measurable P&L impact.



The five pillars of AI risk in 2026

1. The shifting regulatory landscape

As we move through 2026, the scrutiny on automated decision-making has intensified.

- **The risk:** General-purpose LLMs are non-deterministic and opaque by design. They cannot produce a “watertight” audit trail explaining exactly why a claim was declined or a policy was priced a certain way.
- **The requirement:** Any architecture must be fully auditable and able to demonstrate human oversight at every juncture. If a regulator asks for the logic behind an AI-driven decision, “the model said so” is no longer a legal defense.

2. The legacy integration gap

Outdated core systems remain the primary barrier to scaling AI. While a LLM can read a PDF, an Agent must act on it.

- **The challenge:** Building agents that connect to live Policy Administration Systems (PAS), claims systems, and billing platforms without disrupting operations or endangering privacy and security is a massive engineering hurdle.
- **The reality:** Internal teams frequently underestimate the complexity of data cleaning and orchestration. Only 7% of insurers have successfully scaled AI initiatives across their entire organisation because they cannot bridge the gap between the modern AI stack and the legacy core. (DeMarco, 2026)

3. The ROI imperative and “pilot purgatory”

More than 80% of insurers dedicate at least £3.75m+ annually to AI, yet many finance teams remain unable to tie these investments to hard returns. (Evans, 2026)

- **The hurdle:** Only 19% of executives report full clarity on AI ROI. Many IT projects get stuck in a cycle of endless testing because they lack the domain-specific KPIs to prove value. (Evans, 2026)
- **The deadline:** IT leaders must show a credible path to measurable P&L impact, such as reduced loss ratios or operational expense, within 12 months, or risk losing budget to business units looking for “off-the-shelf” alternatives.



As we move through 2026, the scrutiny on automated decision-making has only intensified.

4. Accountability and the talent gap

When an agent makes a consequential error, such as miscalculating a medical treatment plan or missing a sophisticated fraud signal. Someone must answer for it.

- **The talent shortage:** 46% of leaders identify a skills gap as their primary obstacle (Ellingrud, 2025). You aren’t just competing with other insurers for AI talent; you are competing with Silicon Valley and the rest of the business community. The need for these skills has skyrocketed, and there isn’t enough talent to fill all these jobs in frontier technologies (Ellingrud, 2025).

- **The liability:** API providers explicitly disclaim responsibility for model outputs. If you build in-house, your team is solely responsible for defending a decision process that may be technically impossible to reconstruct after a model update.

5. Model drift and vendor lock-in

Insurance workflows cannot tolerate unpredictable model drift. An overall model drift for any agent can cause operational havoc.

- **The requirement:** Sophisticated architectures must be model agnostic. To future-proof the stack, IT leaders need a “wrapper” that allows them to swap underlying LLMs (from GPT-4 to Claude to Llama) without disrupting the specific insurance workflow layer or the human-in-the-loop triggers.

Specialised agents in action: beyond basic chat

To move beyond simple summarisation, Agentic AI must be specialised for the specific complexities of

insurance. Shift’s architecture utilises a multi-agent system designed for critical work areas:

- 1. Fraud and risk agents:** Agents accurately detect fraudulent claims, documents and policies, guide investigations to increase impact, and accelerate case management.
- 2. Coverage and liability agents:** Quickly review coverage, dissect claims details, and prioritise next best actions to produce accurate liability estimates with defensible logic.
- 3. Personal injury agents:** These assess personal injury and Employers Liability documentation, develop and continuously adapt case progression, and guide handlers to accurate outcomes for the injured party and insurers.
- 4. Subrogation agents:** Identifying recovery opportunities often requires digging through weeks of manual case-building. Agents can identify these opportunities in minutes, create the initial outlay request, and guide the negotiation process.

Visualising the strategic decision: build vs. partner

Critical capability	DIY / in-house build	Shift agentic AI
Training data	Limited to your own historical data; requires massive cleaning.	4 Billion+ records analysed; trained on global insurance patterns.
Domain expertise	Requires hiring specialised AI/insurance hybrid engineers.	200+ dedicated insurance data scientists at your disposal.
Governance	Custom-built audit trails; manual “human-in-the-loop” coding.	100% explainable by design; regulatory-ready audit logs.
Time-to-value	18–24 months for a production-grade environment.	Production deployment in months, not years.
Integration	Manual API building for every legacy system.	Pre-built connectors for major PAS and Claims platforms.



The competitive reality: key figures

“The question isn’t whether your engineering team could build this. They probably could. The question is: at what cost, over what timeline, and what else won’t get built while they do?”

- **6.1x:** AI leaders in insurance have generated 6.1x higher total shareholder return than laggards over the last five years. (Nick Milinkovich, Sid Kamath, Tanguy Catlin, and Violet Chung, with Pranav Jain and Ramzi Elias, 2025)

- **98%:** Shift’s customer retention rate, reflecting the stability of production-scale AI compared to the volatility of internal experiments.
- **10 years:** The amount of time Shift has spent refining insurance-specific AI. That is a decade of “edge case” knowledge that cannot be replicated by a generic prompt.
- **120+:** Global customers currently live in production, moving past the “PoC” phase into actual operational transformation.

Conclusion: the path to 2027

The gap between AI leaders and laggards is no longer just a matter of innovation: it is a matter of survival. The carriers who move from pilots to production fastest will set the pace that everyone else has to match.

The real value of Agentic AI isn’t in the model itself, but in the insurance-specific data, expertise, and regulatory readiness that makes the model workable in a high-stakes environment. Shift provides that ready-made architecture, allowing your IT team to focus on strategic implementation rather than reinventing the foundational wheel.

Ready to see what purpose-built looks like?

Contact us to walk through how our agent architecture handles your critical processes.

About Shift Technology

Shift delivers AI agents that transform insurers’ most critical work. By combining deep industry expertise with unrivalled data resources, Shift provides proven results that have earned the trust of hundreds of the world’s leading insurers. Our insurance-grade AI is accurate, explainable and secure, empowering human experts to act with unparalleled speed, complete confidence and a renewed focus on the people they serve.

Learn more at www.shift-technology.com/en-gb

SHIFT