

AI Adoption in Swiss Health & Life Claims Operations

Market Insights Report 2026

Qantev x Swiss InsurTech Hub

01 Executive Summary

Swiss health and life insurers are no longer asking *whether* to use AI; they are asking how to make it pay off at scale. Our interviews with leaders from the largest Swiss H&L insurers give a clear and consistent message: **building AI is the easy part. Deploying it at scale is the real work.**

KEY FINDINGS

- **Governance comes first.** The leading insurers keep the "PoC graveyard" small by filtering use cases early. They use central approval, strategic project portfolios and "lighthouse" pilots that must prove their value before rollout.
- **Integration is the main barrier.** Most of the effort goes into connecting AI to legacy core systems and processes, not into building models.
- **People are the bottleneck.** The same scarce IT and data teams build pilots, run solutions in production and maintain them. As portfolios grow past ten projects in parallel, this split becomes the limiting factor.
- **Claims results are real and measurable.** Efficiency targets are already being met ahead of plan, and specialized models have sharply raised document intake accuracy.
- **The next frontier is agentic AI.** Insurers are moving from AI assistants to autonomous agents, aiming at full end-to-end automation of simple claims.

70–80%

of AI project effort goes into integration with core systems

30–40%

of AI pilots reach enterprise scale

15%

claims efficiency gain, reached ahead of target

98%

intent recognition on invoices with specialized models (up from ~72%)

02 Context & Methodology

As health insurance premiums in Switzerland keep rising, operating margins are under pressure, and insurers are turning to AI to streamline claims and contain costs. Global research shows a clear gap, though. McKinsey reports that **88%** of enterprises use AI in at least one business function, but only **1 in 3** have scaled beyond pilots. BCG finds that **74%** of companies have yet to show tangible value from AI, and only **4%** consistently generate significant value. BCG also traces only **10%** of AI challenges to the algorithms themselves; around **70%** stem from people and processes.

This study looks at whether that gap exists in the Swiss H&L market, and why.

HOW WE RAN THE STUDY

- **Participants:** interviews with 5+ senior data, AI and IT leaders at the largest Swiss health and life insurers.
- **Format:** 30-minute structured interviews held between July and October 2026, all following the same three-block framework (Strategy · Processes & Challenges · Results & Outlook) so answers can be compared.
- **Confidentiality:** all answers are anonymized. No company names, individuals or proprietary data are disclosed.
- **Reading this report:** findings come from the interviews. Boxes marked "Qantev perspective" give Qantev's own interpretation.

03 Strategy: How Insurers Choose Their AI Use Cases

AI has moved from innovation labs to the board agenda. None of the insurers we interviewed treat AI as a standalone initiative. Instead, AI is built into existing corporate and departmental strategies and managed through the same project portfolio processes as any other IT investment.

We saw three governance models:

Centralized steering

A central AI team approves every proof of concept before it starts. Use cases are ranked by strategic fit (e.g. cost efficiency in a turnaround phase, or growth in new segments), by the number of people working in the process, and by the available IT capacity.

Federated lighthouse projects

Each large business area (claims and products, client and market, IT) runs its own flagship pilot. A cross-functional AI decision body meets monthly, and pilots that prove their value are rolled out to other departments on a shared, standardized platform.

Three-speed streams

AI is managed in three streams that move at different speeds: people, process and technology. AI ambassadors in each department surface use cases from their teams and share information back.

In all models, ROI is measured in operational terms: claim volumes, average handling time and time saved per case. Participants put AI at around **10–20% of the IT project budget**. Typical timelines are **6–9 months for a pilot** and roughly **one more year for enterprise rollout**, with roadmaps running into 2027–2028. Some insurers plan every AI business case over three years at most, assuming the solution may be obsolete after that.

04 From Pilot to Production: The Scaling Barriers

This is the core of the study. Across interviews, participants agreed that the difficulty is not in building AI but in *industrializing* it. As one put it: “Building a first PoC is quite simple and quite fast. The huge effort is the industrialization step.”

Barrier 1 – Integration with legacy core systems. No Swiss insurer runs a fully modern, API-ready core system. Scaling an AI solution means real work on data integration, API connections, identity and access management, and embedding the solution into the claims handler's daily workflow (for example, a single button in the core claims system that opens the AI assistant with the right case context).

Barrier 2 – Limited people and IT capacity. Those who build pilots are usually the same people who must then run and maintain them. This works for a handful of projects, but it breaks down once the portfolio grows. Core-system replacement programs and other IT priorities compete for the same teams, so insurers must prioritize by *IT capacity*, not only by expected ROI.

Barrier 3 – Change management. Moving from pilot to daily operations also requires monitoring, training and adoption by end users. Participants see this as a separate phase in its own right, not an afterthought.

Barrier 4 – Regulatory uncertainty. With future AI rules still unclear, use cases close to the line are postponed. Insurers manage the risk by prioritizing low-risk use cases first.

Does the “PoC graveyard” apply in Switzerland? Only partly. Some participants estimate that **30–40% of AI pilots reach enterprise scale**. They credit strict filtering before any PoC starts: smaller organizations cannot afford hundreds of experiments, so central approval acts as a natural brake. But filtering works by testing fewer ideas, and even then **60–70% of pilots do not reach scale**. They wait on the backlog rather than being abandoned: a quieter version of the same graveyard. Some insurers run proofs of value (PoVs) with KPIs set upfront instead of technical PoCs.

“ Building the AI models is no more than 20–30% of the effort. 70–80% is really the integration into your core systems and core processes.

20–30% 70–80%

- ☒ Building the AI models
- ☐ Integration into core systems & processes

QANTEV PERSPECTIVE

Any insurer can now call a foundation model. What isn't a commodity is the **harness** around it: claims rules, medical coding logic, confidence thresholds, workflow integration and monitoring. That harness is the 70–80%, and every insurer that builds it alone rebuilds the same thing with the same scarce team.

“ The people working on the pilots are nearly the same ones who have to operate and maintain the solution. With two or three projects that's fine, but now we're talking about more than ten.

05 Impact on Claims Operations

Claims is where AI already shows its clearest, measurable impact. The use cases in production today fall into three groups:

Document intake

Up to 98%

intent recognition on incoming invoices, after some insurers replaced generic, off-the-shelf models (~72%) with specialized models trained on their own claims documents.

Automated invoice checks

70–80%

of simple health invoices are processed automatically. Structured e-bills from providers are the easiest to automate; paper bills sent in by patients are the next target.

Claims assistants

Human in the loop

AI copilots summarize long claim files, draft emails for missing information, pre-check coverage and suggest the next best action, while the handler confirms every result.

“ Our target was a 15% efficiency increase in claims by end of 2026. We already achieved it, half a year early.

Participants still name one main bottleneck: **how resources are split between maintaining solutions already live and delivering new ones.**

Not every insurer sees claims as its biggest AI win. Some had already automated about **80% of claims** before AI, so the larger gains came from customer support, internal operations and developer productivity. Foreign bills stay manual for now, because of their variety and higher fraud risk.

QANTEV PERSPECTIVE

The jump from ~72% to 98% is a lesson in **specialization**. It matters most in what is left to automate: paper bills, foreign invoices and complex cases. That is where rules engines stop and generic models are least reliable, with unfamiliar formats, many languages and coding systems, and higher fraud risk.

06 Outlook: Toward Agentic, End-to-End Claims

The next step is a move from *assisted* to *autonomous* operations. Today, AI supports the claims handler, who still copies the drafted email, confirms the coverage check and decides the next step. Tomorrow, AI agents will carry out these steps on their own for simple cases.

Participants expect agentic AI to roughly **double today's efficiency gains (up to ~30%)**, though these figures remain an ambition and are not yet in the business plan. Early agentic PoCs are running in 2026, with production targeted from 2027.

Beyond back-office automation, insurers also plan to use conversational AI (chatbots and voice bots) to answer customers' claims questions directly.

“ Full end-to-end automation of simple claims, from first notice of loss to settlement, is where we see the next level.

QANTEV PERSPECTIVE

Autonomy is earned one decision at a time: an agent can settle a claim alone only if each step is scored, explained and reversible. **Today**, a low confidence score sends the case to a handler. **Tomorrow**, high scores let simple claims run end to end.

07 Key Takeaways

1 Budget for integration, not just models. Expect most of the effort and cost to go into connecting AI to core systems and processes.

2 Filter early and centrally. Strict governance before a PoC starts is the most effective defense against the PoC graveyard.

3 Plan the people side from day one. Separate or grow the teams that build, run and maintain AI, or the portfolio will stall as it grows.

4 Start in claims, then go agentic. Document intake, invoice checks and claims assistants deliver measurable ROI today. End-to-end automation of simple claims is the next wave.

08 How Qantev Can Help

Qantev is the Claims Intelligence Platform for health and life insurers. Generic AI is good at language: summaries and drafted emails. Claims decisions need more: a harness of claims rules, medical coding logic and confidence thresholds that carries the reasoning of claims operators, medical doctors and fraud specialists. Qantev checks every claim before payment, so clean claims no longer wait behind complex ones, and your experts decide what is automated.

1 billion

medical claims processed

20+

clients in 10+ countries

Up to 90%

faster end-to-end turnaround, from claim receipt to decision

< 12 months

to positive ROI for every client, well inside a three-year business case

Claim Intake

Classifies documents, extracts data and predicts medical codes in any language. Each extraction is scored, so you choose full automation or manual review.

Claim Insight

Automated coverage checks, quote analysis and medical-coding coherence checks for complex claims and pre-authorizations.

Fraud, Waste & Abuse

Flags suspect claims and providers before and after payment, using medical anomaly detection and network analysis. Every alert comes with its reasons.

HOW WE ANSWER WHAT THIS REPORT FOUND

- **Integration & IT capacity.** No integration needed to start, so the use case doesn't wait in the IT queue. Flexible APIs then connect Qantev to your core claims system.
- **Specialization & control.** Models trained on your own claims, with a confidence score on every decision. Your experts decide what is automated.
- **Model churn.** A platform, not a one-off build. Keeping models current is our job, so a three-year business case doesn't end in a rebuild.
- **People capacity.** A dedicated Qantev team from discovery to go-live and beyond: we run, monitor and retrain models in production, freeing your team for the next use case.
- **Regulation.** Auditable reasoning on every decision, whatever the final AI rules. ISO 27001 certified, aligned with the Swiss FADP, with Swiss data hosted in Switzerland.
- **The hard tail.** Built for what rules can't handle: paper and foreign bills across languages and coding systems, drawing on claims experience in 10+ countries.

SKIP THE POC GRAVEYARD

A Proof of Value in 4 weeks or less, on your own data

Pilots in this study typically take 6–9 months. Ours sets the KPIs upfront and needs no integration.

1

Scope

Pick the use case and agree the success KPIs.

2

Share data

A representative sample of historical claims. No system integration.

3

Run

Qantev runs the solution on your claims in its secure SaaS environment, hosted in Switzerland.

4

Decide

Results against the agreed KPIs, and a business case built on your own numbers.

Ready to move from pilot to production?

Trusted by insurers and TPAs including **AXA, Bupa, Generali** and **Global Excel**.

contact@qantev.com

www.qantev.com

This report is based on anonymized interviews with leaders of Swiss health and life insurers. Sources for global benchmarks: McKinsey & Company, "The State of AI in 2025: Agents, Innovation, and Transformation"; Boston Consulting Group, "Where's the Value in AI?" (October 2024).