



The state of automation in Insurance, 2026

From experimentation
to execution

Executive summary

Insurance has crossed a quiet but decisive threshold.

For most of the last decade, AI lived at the edges of the business—pilots, proofs of concept, innovation labs that demonstrated promise but rarely reshaped how work actually got done. That era is over. By late 2025, the majority of insurers had deployed generative AI somewhere in production. The question executives are asking now is no longer whether AI belongs in insurance, but how to turn it into durable operational advantage.

That shift in mindset matters. It reflects a broader realization taking hold across carriers, brokers, managing general agents (MGAs), and reinsurers: incremental automation and standalone AI tools are not enough. The real opportunity lies in rethinking end-to-end journeys—customer, broker, and operational—and enabling AI to do more than advise. It must act.

This is where agentic AI enters the conversation.

Agentic AI represents an evolution in AI. It's AI that executes within defined guardrails, across systems, and under human oversight. Done well, agentic AI allows insurers to take work off people's plates: reviewing submissions, triaging claims, orchestrating document gathering, managing routine service interactions, and coordinating complex, multi-step processes that have historically resisted automation.

In this paper, I argue three things:

First, agentic AI is not "just another chatbot." It is the mechanism by which insurers finally close the gap between AI insight and operational action.

Second, the biggest gains come from applying agentic AI to journeys, not isolated use cases—specifically the policyholder journey, the broker journey, and the deeply manual or outsourced processes that sit behind the scenes.

Third, speed to value, governance, and architectural flexibility are no longer optional. Insurers need results in months, not years, and they need AI systems that are observable, compliant, and resilient as models, regulations, and business priorities evolve.

The goal of this white paper is not to sell a vision of AI's future. It is to offer executives a practical way to think about how agentic AI is already reshaping insurance work—and how to move from experimentation to execution in 2026.

2025 marked the end of "AI as an experiment"

If 2023 and 2024 were about discovery, 2025 was about acceptance.

Across lines of business and geographies, insurers reached a shared conclusion:

AI is no longer a sideshow. It is becoming one of the defining forces shaping cost structures, customer expectations, and competitive differentiation over the next three to five years.

Several forces converged to make this unavoidable.

Generative AI moved from novelty to necessity. Boardrooms and executive teams stopped debating if AI would matter and began debating where it would matter first. In many organizations, generative AI became visible to customers and employees through claims support, contact centers, underwriting assistance, and document-heavy operations.

At the same time, the economic environment left little room for delay. Persistent catastrophe losses, inflation, and competitive pressure continued to squeeze margins.

Customers and brokers expected faster, clearer, more digital experiences—without corresponding increases in headcount. Simply adding people was no longer a viable lever.

By the end of 2025, most insurers had tangible AI activity underway. What varied was how far that activity actually went.

Many organizations accumulated a portfolio of pilots: document summarization tools, first notice of loss (FNOL) triage experiments, chatbots for service teams, early underwriting co-pilots. Some teams achieved real, localized wins. But at scale, a familiar pattern emerged.

AI generated insight—but people and traditional automation still had to do most of the work.

Fragmented stacks made it difficult to govern, monitor, or extend solutions. Risk and compliance teams raised concerns about explainability and control. And too often, promising pilots stalled before becoming core capabilities.

The issue was not that AI failed. It was that AI, deployed in isolation from automation and operating model change, hit a ceiling.

What we actually mean by agentic AI in insurance

“Agentic AI” is quickly becoming one of those phrases that is widely used and rarely defined. For insurance leaders, clarity matters—because the distinction between agentic AI and earlier AI approaches is not semantic. It is operational.

In practical terms, agentic AI refers to AI software agents that can perceive, reason, and act on behalf of people or systems across multiple applications and data sources, within clear guardrails and with humans in the loop where required.

An AI agent does not simply answer a question. It works toward a goal.

Given an objective—such as validating a broker submission, triaging a claim, or processing a service request—an AI agent plans a sequence of steps, invokes the appropriate tools along the way, and takes action in systems of record. It reads documents and emails, calls APIs, updates core systems, drafts communications, and knows when to pause and escalate to a person.

This is a meaningful departure from chatbots and standalone machine learning models. The difference is not intelligence; it is ownership. AI agents take responsibility for defined tasks and move work forward, rather than stopping at recommendation.

This shift is emerging now for a simple reason: the prerequisites are finally in place. Language models are capable enough to reason over complex documents and unstructured inputs.

Automation and integration layers are mature enough to allow reliable action across legacy systems. And business pressure has intensified to the point where incremental gains are no longer sufficient.



Where agentic AI is already delivering value

Insurers do not need another catalog of AI use cases. They need clarity on where agentic AI is already producing measurable outcomes.

Across carriers and brokers, three journeys consistently stand out.

The policyholder journey: from quote to claim

This is where much early GenAI experimentation began—and where agentic AI is now starting to do real work.

Today, friction is common. Quote intake arrives through multiple channels. FNOL submissions are often incomplete. Claims teams spend hours assembling a coherent picture from documents, photos, emails, and third-party reports before meaningful decisions can be made.

An agentic approach changes the shape of that work.

Intake AI agents normalize submissions from web forms, emails, and attachments into a single structured view. They prepopulate systems, flag missing information, and prepare draft quotes for review. Claims AI agents ingest FNOL from any channel, classify complexity, gather evidence, and propose next actions—while keeping policyholders informed with clear, empathetic updates.

The impact is not theoretical. Insurers applying agentic AI to bounded segments of the claims and service journey are seeing shorter cycle times, higher straight-through processing on low-complexity cases, and measurable improvements in customer satisfaction. The difference is that work actually progresses without waiting for manual handoffs at every step.

The broker journey: from submission through renewal

For insurers that sell through intermediaries, the broker experience is the commercial front door. Ease of doing business often outweighs product differences.

Yet brokers still contend with rekeying data, slow responses, and opaque renewal processes. Agentic AI directly addresses this friction.

AI agents monitor inboxes and portals, classify submissions, extract data from Association for Cooperative Operations Research and Development (ACORD) certificates and attachments, and prepare underwriters' first-look summaries. Routine endorsements and service requests can be handled end to end for standard cases, with complex exceptions routed to people alongside clean context and suggested actions.

In renewals, AI agents can begin work weeks in advance—pulling performance data, drafting strategies, and nudging brokers for required information—so underwriters are not scrambling at the last minute.

The result is faster responses, less back-and-forth, and more capacity for underwriters and distribution teams to focus on relationships and higher-value deals.

The manual and outsourced process stack

The third opportunity is less visible to customers but often more consequential financially.

Bordereaux processing, delegated authority management, claims administration, finance close, and regulatory reporting remain deeply manual in many organizations. Outsourcing has helped manage cost, but it has also entrenched document-heavy workflows and reduced transparency.

Agentic AI changes the economics. AI agents can ingest and normalize bordereaux files, flag anomalies, and generate discrepancy reports. They can manage routine follow-ups in claims administration and orchestrate reconciliations across finance systems, producing first-draft reports for people to review.

This is not about eliminating partners overnight. It is about rebasing work so that both internal teams and providers operate in a more automated, observable environment—reducing leakage, improving control, and freeing capacity for higher-value analysis.

The technical reality executives cannot ignore

As agentic AI moves from concept to production, four realities quickly assert themselves.

First, time to value matters. Boards and executive teams expect visible impact within 90 to 180 days on clearly defined journeys. Agentic AI cannot become a multi-year science project.

Second, governance is non-negotiable. Every AI agent action must be observable and auditable. Insurers need clear guardrails, human-in-the-loop controls, and the ability to explain decisions to regulators and internal risk committees.

Third, build-it-yourself orchestration is a trap for most organizations. While it promises control, it often leads to brittle systems, escalating engineering costs, and difficulty keeping pace with rapid model and regulatory change.

Finally, architecture must remain flexible. Models will evolve. Regulatory expectations will shift. Locking the business into a single model or closed ecosystem is a strategic risk.

The insurers that scale successfully treat models as replaceable tools, AI agents as productized components, and orchestration as a governed layer that spans systems, models, and automations.

A pragmatic path forward in 2026

The most important decision is not whether to pursue agentic AI. It is how to start.

Successful organizations do not begin with sweeping transformation programs. They begin with a small number of anchor journeys, clear ownership, and disciplined learning.

In the first months, leaders focus on where friction is highest and volume is sufficient to matter. They establish basic governance and deploy initial agents in production, not in isolation.

Over the following quarters, they codify what works—shared document skills, integration patterns, observability—while expanding deliberately into adjacent journeys. Only later do they formalize operating models and workforce strategies based on real data, not assumptions.

The common thread is focus. Agentic AI delivers value when it is applied to real work, with real accountability, under real constraints.

Closing thoughts

Agentic AI is not about chasing the latest demo. It is about changing how work gets done in insurance.

It reduces friction in customer and broker journeys. It pulls cost and risk out of manual and outsourced processes. And it gives people leverage—so they can focus on judgment, relationships, and innovation instead of repetitive coordination.

The technology is ready enough. The question now is organizational: who will design the journeys, own the agents, and build the governance that makes this sustainable?

The insurers that answer that question in 2026 will not just have better AI. They will have better businesses.

Key contact



Eric Stroom
Insurance Practice Leader, UiPath

eric.stroom@uipath.com