



Faheem Ballim

Senior Consultant
Digital Consulting

Tel: +27 60 997 6473
Email: faheem.ballim@kpmg.co.za



Tasneem Essop

Consultant
Digital Consulting

Tel: +27 64 758 2647
Email: tasneem.essop@kpmg.co.za

The evolution of insurance intelligence

The insurance industry is founded on the promise of protecting people and businesses from uncertainty. Insurers ingest immense volumes of data, assess risks and pay claims, while satisfying regulatory requirements and employing measures in controlling costs. Artificial intelligence (AI) has already delivered incremental improvements to insurers using predictive analytics and basic automation. Now, a new chapter is emerging with **agentic AI** - a class of intelligent systems that can autonomously reason, plan and act. This article traces the progression from rule-based automation to AI agents and finally to agentic AI, explores recent adoption trends and statistics, deep dives into delegated authority processes like bordereaux and binder management, highlights wider insurance applications and outlines practical considerations for deploying these technologies responsibly.

From automation to agency: evolution of AI systems

Rule-based automation and robotic process automation

Before the rise of machine learning, insurers relied on **rule-based automation**, often labelled **robotic process automation (RPA)**, to handle repetitive back-office tasks. These bots follow explicit scripts and manipulate structured data, performing deterministic actions such as policy issuance or invoice generation. RPA excels when inputs are consistent and decision rules are unambiguous; it cannot interpret unstructured data or adapt to new situations without re-coding. Analysts note that RPA is ideal for repetitive, rules-driven processes, whereas more intelligent agents are needed when tasks involve unstructured information or dynamic decision-making. Each regulatory change or product update requires developers to adjust scripts, illustrating the fragility of rule-based approaches.

AI agents: autonomous and adaptable

The next stage of machine learning introduced **AI agents**, which perceive their environment, reason, make decisions and take action to achieve predefined goals with limited human intervention. Insurers have adopted chatbots that answer customer queries, underwriting assistants that assemble risk summaries and claims bots that pre-screen documents. Unlike RPA, these agents incorporate machine-learning models and generative AI (Gen AI) models, enabling them to interpret text, images and speech and to learn from data and humans in the loop feedback.

They can be **reactive** (responding to user requests), **proactive** (initiating actions based on triggers), **goal-oriented** (planning steps towards outcomes) or **hybrid** (combining behaviours). Adoption is broad: health insurers use AI agents for customer service, while property insurers deploy them for claims triage and fraud detection. Over time these agents improve performance by learning from historical outcomes.

Agentic AI: reasoning, judgment and collaboration


Moving on to recent years, **agentic AI** is the latest wave in machine learning. It combines Gen AI models with reasoning and planning capabilities so that systems can decompose complex tasks, set their own goals, invoke specialised tools, monitor progress and collaborate with humans or other agents. In insurance, an agentic AI system can ingest a claim, verify coverage, request missing information, compute settlement options, draft customer communications and update the policy in one coordinated workflow which is able to handle various nuances. This capability, coupled with **retrieval-augmented generation (RAG)**, ensures that generative models consult the right documents before providing an answer, increasing accuracy and context of execution and delivery. Agentic AI and AI agents are emerging as a top investment priority: a recent KPMG study revealed that almost **three-quarters of insurance CEOs** view generative AI as their most important opportunity. Market analysts estimate that AI investment across the global insurance sector could reach **USD79 billion by 2033**, reflecting expectations for agentic systems to transform underwriting, claims and distribution¹.

¹ KPMG, "KPMG Q1 2025 AI Pulse Survey," 2025.

From RPA and AI agents to agentic AI: key differences

Progressing from RPA to AI agents and then to agentic AI reflects a shift from scripted automation to systems that understand context, learn from data and collaborate. **RPA** automates predictable, high-volume tasks; **AI agents** interpret unstructured data and make independent decisions to solve a specific workflow/use-case; **agentic AI** plans multistep goals, co-ordinates multiple agents, orchestrates other tools based on discretion and adjusts actions based on feedback and through self-set goals. Analysts recommend using RPA for deterministic workflows and deploying AI agents when tasks involve interpretation or adaptation for an entire workflow. For example, an AI agent can analyse customer sentiment or summarise free-form text capabilities, of which RPA is not capable. Agentic AI goes further: by combining generative models with planning and dynamic self-set goals, it decomposes complex goals, calls on calculators or search engines, monitors outcomes and reflects on performance using its own discretion. Compared to the use of scripted mechanisms applied with RPA, agentic systems are more versatile but require guardrails and greater computational resources. Insurers are encouraged to adopt a continuum of these types of technologies, leveraging RPA for stable processes while deploying AI agents or agentic AI where adaptability and learning are essential.

Figure 1: Definition of intelligent automation, agents and agentic AI



Market momentum: investment, adoption and governance

Investment momentum and return on investment (ROI)


Investment in AI is on the rise. KPMG's **Q1 2025 AI Pulse Survey** reports that organisations plan to invest an average of **USD114 million** in generative AI over the next 12 months, up from **USD89 million** in Q4 2024. Executives cite profitability and productivity as the key benefits of Gen AI deployment: **97%** say that investment in Gen AI drives profitable returns and **94%** highlight productivity gains. In Q2 2025, budgets are being allocated to foundational safeguards with **67%** of Gen AI budgets going to cyber and data security, **52%** to risk and compliance and **48%** to operational activities. These results indicate that insurers are moving away from experimentation and starting to focus efforts on operational integration and prioritising secure deployment^{1 2}.

Adoption patterns and sourcing strategies

Survey data reveal a rapid shift from exploration to piloting and deployment. In Q1 2025, pilot programmes for AI agents leapt to **65%** from **37%** in previous quarters, yet full deployment remained at **11%**. By Q2 2025, **90%** of organisations were beyond experimentation: **33%** had deployed agents, **57%** were piloting and **10%** remained in the exploration phase.

Sourcing strategies also evolved, in Q1 2025, two-thirds of leaders (**67%**) planned to **buy** pre-built agents, **27%** intended to combine buying and building, with **5%** planned to build internally. Moving into Q2 2025, **51%** of organisations planned investment in a combination of pre-built and internally built agents. Those organisation planning to buy dropped to **46%** and those planning to build dropped to **2%**.

Importantly, **93%** of Q1 2025 respondents said their investment in Gen AI had already **enhanced their competitive position**¹.



² KPMG, "AI Quarterly Pulse Survey: Q2 2025," 2025


Drivers and challenges in AI adoption

The uptake of artificial intelligence across industries is shaped by a complex interplay of **drivers that accelerate adoption** and **challenges that constrain it**. Understanding both dimensions is essential for organisations seeking to embed AI into core operations.

On the **demand side**, three drivers stand out:

1. The pursuit of **efficiency and cost optimisation** continues to dominate executive agendas. Gen AI in particular is viewed as a tool to automate labour-intensive processes, accelerate document review and reduce operational bottlenecks.
2. Organisations are motivated by **innovation and competitiveness**. Executives increasingly recognise that AI can create new revenue streams, enhance customer engagement through personalisation and support differentiated product offerings. In an environment where first-mover advantage is significant, the risk of inaction is becoming as material as the risk of adoption.
3. **Technological maturity** has lowered barriers to entry. The scalability of cloud platforms, availability of foundation models and emergence of specialised vendor ecosystems have reduced both cost and complexity, enabling firms to move more confidently from pilot projects to enterprise-scale deployments.

Against these drivers are a set of persistent **challenges**. Foremost among them is **risk management**. In KPMG's Q1 2025 Pulse survey, **82%** of executives cited risk as the primary obstacle to implementing their AI strategies. Concerns span model reliability, data leakage, ethical use and reputational exposure. Closely related are anxieties around **data quality (64%)** and **trust in AI (35%)**, underscoring the dependence of AI systems on reliable inputs and the importance of transparency in decision-making¹.



The workforce dimension adds another layer of complexity. The Q1 2025 Pulse survey identified **system complexity (66%)**, the **pace of technological change (56%)** and **technical skills shortages (51%)** as the most pressing barriers to adoption. By Q2 2025, these concerns had shifted but remained acute: **skills gaps were cited by 59% of executives**, followed by **workforce resistance to change (47%)** and **system complexity (39%)**. Together, these findings point to a widening gap between the rapid evolution of AI technologies and the availability of talent capable of deploying, managing and governing them. The cultural component is equally significant: without clear communication of benefits and alignment to business strategy, employees may resist or underutilise AI tools¹.

Finally, **regulatory and governance concerns** have moved to the forefront. In Q2 2025, **data privacy (69%)**, **regulatory compliance (55%)** and **data quality (56%)** emerged as leading issues. As jurisdictions worldwide accelerate the development of AI legislation, organisations face increasing pressure to demonstrate compliance, safeguard sensitive information and establish robust oversight frameworks¹.

Taken together, these findings highlight a dual reality: while AI presents a powerful engine for efficiency and innovation, its potential can only be realised through parallel investments in **governance, data quality and workforce capability**. Organisations that succeed will be those that not only seize the opportunities but also systematically address the risks that accompany them.

Insurance value from agentic AI and RPA: adoption, workforce and governance considerations


In insurance, the case for agentic AI and RPA is fundamentally operational: shorten cycle times, reduce leakage and lift decision quality at scale. Agentic AI and AI agents (autonomous or semi-autonomous “agents”) complement RPA’s deterministic workflow automation by handling ambiguous, language-heavy tasks, such as reading documents, conversing with customers and coordinating next steps across systems. RPA executes the repeatable steps reliably and at a low cost. Together they form an orchestration layer across the value chain:

- **Distribution and servicing:** agents act as licensed-assistant copilots for brokers and contact-centre staff, summarising prior interactions, drafting compliant responses and pre-populating quotes and endorsements. RPA posts updates to core policy administration and client relationship management systems and triggers downstream tasks. Result: higher first-contact resolution and faster turnaround times in delivering quotations or signing on clients.

- **Underwriting:** intake agents triage submissions, extract data from broker packs and surface anomalies; RPA validates mandatory fields and pushes clean risks to straight-through processing. Underwriters spend time on risk selection instead of rekeying.
- **Claims (first notification of loss (FNOL) to settlement):** agents guide FNOL conversations, classify coverage questions and assemble evidence; RPA orders reports, reconciles invoices and issues payments once rules are met. Net effect: shorter time-to-decision, less leakage and better claimant experience.
- **Fraud and risk:** agents synthesise unstructured signals (adjuster notes, call logs, social content where lawful) and prompt human review; RPA operationalises watchlists and case routing. This augments the productivity of claims investigation units without overwhelming investigators.
- **Finance and compliance:** report-drafting agents compile bordereaux and regulatory templates from ledgers and data marts; RPA reconciles journals and attestation workflows, reducing friction while preserving control.

Executives see this cross-functional upside clearly: the greatest expected benefits are in **technology (76%)** and **operations (74%)**, with **risk (56%)** following closely; gains also extend to **finance (39%)** and **marketing/sales (35%)** where faster analytics and personalised outreach matter. The adoption pattern mirrors the value thesis: **daily use of AI productivity tools is at 58%**, **weekly use of knowledge assistants (e.g., RAG) at 61%** and **35% of organisations report embedding Gen AI directly into operational workflows**. Critically, leaders frame the impact as augmentative: **76% expect automation of routine tasks without eliminating roles**, **69% anticipate freeing high performers for strategic work** and **57% foresee uplift for struggling performers** - a shift from labour substitution to capability elevation.

With greater autonomy comes tighter guardrails. **Sixty-three percent** plan to deploy agents from trusted providers and **52%** report mechanisms being in place that **prevents access to sensitive data without appropriate human oversight**. This does **not** mean AI independently restricts data. It means organisations are instituting **human-defined governance controls**, such as role-based access, approval checkpoints, redaction and audit trails, within which agents operate. Humans set policy and agents execute inside those boundaries.



Board capability is rising but uneven. Only **8%** of executives report substantial AI expertise at board level, while **69%** note moderate proficiency; nonetheless **45% of boards discuss AI at every meeting** and **41% do so frequently**. For insurers, the next step is formalising oversight: articulate the AI risk appetite, link model controls to conduct and prudential obligations and tie investment cases to combined-ratio and customer-outcome metrics.

Bottom line for insurers: the simultaneous application of agentic AI and RPA is already translating into throughput, accuracy and experience gains across underwriting and claims, with workforce adoption advancing and governance processes maturing in parallel. Successful organisations will be those who approach value realisation, workforce enablement and governance as an integrated operating model – designed, scaled and validated together¹.

Figure 2: Benefits across intelligent automation and agentic AI



Efficiency gains and impact metrics

Evidence from field deployments indicate that AI can deliver material improvements across insurance operations. Consulting research and analysis of AI transformations in property and casualty underwriting and distribution found a **10-20%** increase in new agent success and conversion rates, **10-15%** premium growth, **20-40%** reduction in onboarding costs and **3-5%** improvement in claims accuracy. A personal lines claims transformation program deployed by Aviva achieved a **23-day** reduction in liability assessment time, **30%** improvement in claim routing, **65%** reduction in customer complaints and more than **GBP60 million** in savings. These results illustrate that the benefits of AI go beyond labour savings: they translate into faster cycle times, improved customer experience, reduced leakage and better decision quality. The next sections consider how agentic AI can apply these lessons to delegated authority and other insurance processes^{1 3}.

Deep dive of a use case: delegated authority workflows

Delegated authority is a core feature of many insurance markets: intermediaries and underwriting management agencies (UMAs) bind policies on behalf of insurers, producing **bordereaux reports** that summarise premiums, claims and risk exposures. Insurers then reconcile these reports against binder terms and sub-ledgers.

While this process is fundamental to distributing business, it remains labour-intensive. Some of these challenges include bordereaux arriving in diverse formats and are often incomplete; binding authority agreements differing across markets; and the performance of manual reconciliations against policies, claims, commissions and taxes. Industry analyses show that manual bordereaux management results in financial and management reporting delays, errors and long standing and high levels of unallocated cash. Intermediaries and UMAs may interpret data fields differently, complicating validation and compliance. Even determining whether a claim falls within a binder's scope may require an underwriter to read the policy and endorsements. Without standardisation, there is a risk of inaccurate entries, omitted or partially captured items (missed exposure accumulations) and regulatory breaches.

³ KPMG, "Advancing AI across insurance," 2024.

Current solutions

Several technology providers have developed solutions: platforms that automate bordereaux ingestion, validation and mapping to a single, standard data template (canonical data model); tools that generate exception reports and queries; and consolidated binder platforms that centralise underwriting, claims and financial data. These solutions can reduce financial reporting close out times and improve data quality. However, they remain largely rule-driven and require manual oversight when exceptions arise.

Agentic AI for bordereaux and binder management

Agentic AI promises to take delegated authority management further by chaining together multiple specialised AI agents. A hypothetical **multi-agent scenario** could work as follows:

- 1 Ingestion and schema (data blueprint) inference:** an intake agent automatically detects new bordereaux files across channels (SFTP, SharePoint, email), infers schema and lineage (the file's structure and the trail of where each field came from) and flags missing or unfamiliar fields. When novel columns appear, it prompts the cover holder or human operator for mapping examples. Over time the agent learns common variations, reducing the extent of manual intervention.
- 2 Normalisation and enrichment:** a transformation agent maps fields to the insurer's canonical data model, enriches rows with underwriting data (territories, perils, currencies) and tags records with binder references. It can incorporate geocodes and hazard codes to facilitate catastrophe accumulation checks.
- 3 Multi-ledger controls:** a control agent reconciles premiums and claims bordereaux with general and sub-ledgers, verifies commission and tax calculations and compares net/gross splits with binder terms. It applies rules for limits, deductibles, sub-limits and exclusions, referencing clauses through retrieval-augmented generation.
- 4 Accumulation and sanctions:** a risk agent updates catastrophe accumulation, monitors aggregation limits, runs sanctions and politically exposed person (PEP) checks and flags suspicious patterns. When rules cannot determine a course of action, the agent retrieves relevant clauses for a human underwriter to review.

- 5 Invoice and settlement:** a finance agent reconciles bordereaux with bank statements and calculates settlement amounts, automatically generating debit/credit notes and identifying unallocated cash. It drafts settlement instructions for sign-off.
- 6 Correspondence and audit:** a communications agent prepares queries to cover holders, attaches exception reports and schedules follow-ups. Every action is timestamped and linked to source data for audit and conduct risk reporting.

By orchestrating these agents, insurers can achieve continuous assurance, reduce manual queries and report exceptions with explanatory context. Human experts remain in the loop, making final decisions and ensuring ethical and regulatory standards are met.

Broader applications across the insurance value chain

This section considers how similar agentic patterns discussed above can be applied across other insurance domains.

Underwriting and product development

AI agents can assist underwriters by summarising submission documents, analysing third-party data and recommending risk scores. They can generate comparative “what-changed” reports when broker submissions are updated and suggest endorsements or coverage terms. Gen AI can help design products tailored to niche markets by analysing customer sentiment, emerging risks and competitor offerings. Future agentic systems could autonomously assemble quote packages and iterate product features based on feedback.

Claims processing and fraud detection

Claims operations can benefit from agents that automate FNOL intake, classify claims and route them to claims assessors based on complexity. Computer vision models can assess damage from photos, while natural-language models can extract information from adjuster notes. Fraud-detection agents can be used for anomaly detection and network analysis to flag suspicious activity, as well as collaborate with human investigators to refine rules. Agentic systems can orchestrate the entire claims journey, prompting claimants for missing information, coordinating with repair networks and recommending settlement options, while ensuring fairness and compliance.

Customer engagement, marketing and distribution

AI chatbots and virtual assistants can be used to handle routine customer queries 24/7, freeing staff for complex interactions. Marketing agents can analyse customer data to target campaigns and personalise offers. In distribution, agents can triage submissions, check appetite against underwriting guidelines and generate quotes. They can also support with automating document generation and compliance checks during onboarding and renewal.

Risk management, compliance and audit

Risk managers can make use of agents to monitor exposures, update catastrophe models and run stress tests. Compliance agents can then cross-check policy wording with regulatory requirements, generate compliance reports and conduct gap analyses. In delegated authority, agents can be used to ensure that required fields are present before submission. Agents can also assist with solvency capital, IFRS 17 and other reporting regimes by extracting data and performing calculations. Audit logs generated by these agents can provide transparency, supporting internal and external audits.

Workforce transformation and talent

Agentic AI will reshape the nature of the skillset required in the insurance industry. Routine tasks will be automated, enabling professionals to focus on analytics, advisory tasks and customer relationships. Nevertheless, technical skills gaps remain a barrier: we noted above that **59%** of organisations cite such gaps as an obstacle to AI agent deployment. Insurers are encouraged to invest in upskilling of staff, foster collaboration between business and technology teams and ensure that employees understand how to work alongside AI. Governance frameworks should also emphasise fairness, transparency and human oversight³.



Organisations cite technical skills gaps as an obstacle to AI agent deployment



Implementation considerations and controls

The insurance use cases outlined above translate into a common set of implementation needs. Success depends on four pillars: data, architecture, people and responsible AI and ecosystem collaboration.

Data quality and governance

AI agents are only as effective as the data they consume. Insurers should map critical data sources, assess quality and invest in cleansing and integration of data. In delegated authority, this means standardising bordereaux formats, defining mandatory fields and implementing robust validation rules. Organisations must also ensure compliance with data-protection regulations and maintain audit trails for model-assisted decisions. A comprehensive testing and validation plan is essential to confirm model accuracy, stability and monitoring thresholds over time¹.

Architecture and technology stack

Delivering use cases at scale requires a flexible, modular architecture that supports multi-agent collaboration. Many insurers are adopting an agentic AI mesh - a composable, vendor-agnostic framework where small components can be assembled and reused across journeys. Typical layers include customer engagement, decisioning, integration with core systems and shared data/infrastructure services.

Cloud foundations matter for two reasons. Firstly, teams can scale computing up or down for training (teaching models) and inference (running them in production) without large upfront cost. Secondly, shared services can be reused across use cases: retrieval engines to locate the right documents and data, vector stores to hold embeddings for RAG and workflow/orchestration to coordinate agents and RPA. The reuse of shared services shortens build time, improves reliability and simplifies governance.

Change management and responsible AI

Even strong designs stall without adoption. Change management should sit at the centre of delivery: leaders must set a clear vision, tie AI work to measurable business outcomes and create space for experimentation and feedback. In parallel, **responsible AI** practices, such as fairness, transparency, explainability and human oversight, are critical in fostering and maintaining trust. Regulators expect robust governance in respect of data privacy, cyber security and IT risk; insurers should run regular bias and performance audits, implement guardrails and keep human experts accountable for outcomes.

Collaboration with ecosystem partners

Insurance workflows are interdependent. Transforming delegated authority and adjacent processes requires coordination across insurers, brokers, intermediaries/UMAs, reinsurers and technology vendors. **Industry standards** (e.g., ACORD) and market platforms (e.g., Lloyd's initiatives) enable interoperability and reduce friction. Co-developing playbooks and shared tools avoids duplication and accelerates adoption, while clear regulatory guidance encourages responsible experimentation.

With the right **data foundations, agentic architecture, people-and-governance practices** and **ecosystem collaboration**, insurers can scale the benefits of agentic AI while managing risk.

Conclusion

Charting the future of insurance with agentic AI

The insurance sector stands at a critical juncture. After years of incremental automation, **agentic AI** offers the ability to **reason, plan and act** across complex workflows, enabling insurers to deliver personalised products, expedite claims and strengthen risk management. While investment in AI is accelerating, the extent of adoption is disparate.

To capture this opportunity, insurers must **invest strategically** in data quality, flexible architectures and talent. They should adopt a **domain-based approach** to transformation, targeting high-impact areas such as underwriting, claims and delegated authority. Leadership commitment, change management and responsible AI governance are essential in ensuring that agentic AI improves outcomes without compromising trust or fairness. As the industry embraces this new paradigm, insurers can redefine their role, not just as risk providers or risk managers, but as **trusted advisors** who provide seamless, personalised and transparent protection. Those who act decisively will define the next era of insurance and capture the value that agentic AI has to offer.