McKinsey & Company

AI in Asia: Reimagining insurance operations through agentic AI

This report is a collaborative effort by Abhilash Sridharan, Alex Kimura, Angat Sandhu, Azam Mohammad, Bernhard Kotanko, Jan Henrich, and Sid Kamath, with Kamlesh Asrani, Natasha Bhide, and Ritesh Agarwal, representing views from McKinsey's Insurance and Operations practices.



November 2026

AI in Asia: Reimagining insurance operations through agentic AI

Foreword

AI is top of mind for insurance leaders, and for good reason, as the potential benefits are enticing. It's clear that advances in predictive, generative, and agentic AI have elevated aspirations.

The benefits of AI go beyond isolated use cases and siloed pilots. In our work with insurance clients across Asia–Pacific, we have seen how the insurance operating model itself is at a pivotal point. One bold idea stands out: success may not come from including AI in insurance operations. Rather, real transformation occurs when businesses completely reimagine insurance operations domain by domain and build their future state leveraging advancements in AI.

Our approach is to rewire the entire workflow for each operations domain. We work with insurers to rethink how operations function in a world where AI agents can interpret unstructured information, plan and coordinate work across systems, take action, and apply checks while humans retain judgment and accountability where risk and complexity require it.

The starting point is the profit and loss (P&L) outcome the insurer wants to achieve. From there, we redesign the end-to-end journey, clarify workflows and human-agent roles, and make use of AI agents through a common architecture. This helps insurers move from AI pilot projects to end-to-end domain transformation.

Importantly, we treat operations as a profitability, growth, and outcomes engine, not simply a way to manage cost. This report makes the argument for AI's role in improving loss ratios, issuance, conversion, persistency, and customer experience—alongside productivity gains. This is particularly relevant in several Asia–Pacific next-generation growth markets where a pure cost-reduction case for AI is not necessarily compelling.

This report looks at the forces that are converging to make this moment critical for agentic AI in insurance operations in Asia–Pacific. It considers where agentic AI could create economic value and where rules, predictive models, or conventional automation may be more appropriate. It explores practical AI use cases that can help insurers to reimagine the following domains: distribution operations, underwriting and pricing, policy servicing and administration, claims management, and corporate and support functions. Finally, it draws on McKinsey's Rewired framework to suggest a path forward for AI transformation.

Our hope is that this report broadens the conversation about AI and helps insurers in the region find long-term value from it. Ultimately, the question may change from "Where can we deploy AI?" to "How do we reimagine this journey so that profitability and efficiency improve?"

Abhilash Sridharan
Partner, Mumbai

Alex Kimura
Partner, Singapore

Angat Sandhu
Partner, Sydney

Azam Mohammad
Senior Partner, Singapore

Bernhard Kotanko
Senior Partner, Hong Kong

Jan Henrich
Senior Partner, Tokyo

Sid Kamath
Partner, Charlotte

Table of contents

Foreword 3

Executive summary 5

Introduction: From AI pilots to enterprise intelligence 8

Chapter 1: The AI imperative for insurance in Asia–Pacific 10

- Customer expectations in financial services have reset, especially in Asia–Pacific 10
- Margins are under sustained pressure across most insurance lines 12
- Technology is ready for wider use across the business 13
- The opportunity to gain an early advantage is shrinking 14

Chapter 2: High-impact agentic AI use cases to rewire domains 16

- Domain 1: Distribution operations 18
- Domain 2: Underwriting and pricing 20
- Domain 3: Policy servicing and administration 23
- Domain 4: Claims management 25
- Domain 5: Corporate and support functions 27

Chapter 3: The six dimensions of the Rewired framework for at-scale impact 28

- 1. Strategy 30
- 2. Talent and organization 30
- 3. Operating model and governance 30
- 4. Technology 30
- 5. Data 31
- 6. Adoption and scaling with new ways of working 31

How to get started 32

Executive summary

AI is changing how insurers organize and run their operations. Together, predictive, generative, and agentic AI offer a practical way to industrialize the most expensive and judgment-intensive parts of insurance operations. But as AI capabilities become widely accessible, it's likely that advantage could come less from access to the technology and more from how insurers redesign their operating models, make AI part of the day-to-day work, and build the skills and systems to scale it across the organization.

Asia–Pacific insurers in next-generation growth markets often contend with greater operational complexity than their counterparts in other parts of the world. For example, they operate across multiple languages, legacy systems, distribution channels, product generations, and regulatory regimes, sometimes within a single organization. This complexity makes agentic AI potentially valuable in coordinating people, processes, and technology across the enterprise.

Several developments are making agentic AI increasingly relevant to insurance operations in the region:

Customer expectations in financial services have changed. The customer-experience bar is now set outside the industry. Customers in the region compare their insurance experience with apps, digital banks, e-commerce, and instant payments that offer instant authentication, real-time status updates, and conversational service in the local language. By contrast, insurance customers are often expected to read 70-page documents and wait 30 days for a property and casualty (P&C) claim. The next reset is already visible. Customers have moved from being passive recipients of insurance—found by an agent, sold a product, and served on the strength of a personal relationship—to active choosers who research online and pick channels to suit their wealth and sophistication. The customer journey is shifting from a vertically integrated model, in which a customer deals mainly with a single channel from discovery to claims, to a horizontal digital–physical hybrid in which customers move between channels to optimize for transparency, quality of advice, speed, and convenience.

Margins are under sustained pressure across most insurance lines. According to our analysis, between 2018 and 2024, combined operating ratios for health insurance failed to improve in half of the region's ten largest markets, for P&C insurance, combined operating ratios failed to improve in half, and after-tax return on equity in life insurance failed to improve in two-thirds—a picture of stalled profitability rather than a cyclical downturn. This lack of improvement across pockets of Asia–Pacific calls for a different cost model, one in which a meaningful share of judgment-heavy operational work is handled at the marginal cost of computing, with humans focusing on the cases that need their judgment.

Technology is ready for wider use across the business. Gen AI has largely been used for predictive models that work well with structured inputs, repeatable workflows, and narrow decisions. Typically, these models have been less effective when work is unstructured, judgment heavy, or conversational, as is a great deal of insurance operations work. This is where agentic AI excels.

The opportunity to gain an early advantage is shrinking. Several Asia–Pacific insurers are demonstrating what AI looks like at scale. Policies are automatically underwritten in seconds, automated case review ratios are rising, and customers receive claims payments within a day. Those that stall AI programs risk falling behind insurers that can scale at double or triple their pace, and at roughly half the unit cost.

A typical multiline carrier runs about 300 core processes across personal, commercial, health, and life insurance, as well as corporate and support functions, spanning distribution operations, underwriting, policy administration, claims, finance, HR, compliance, and IT. The temptation, given that breadth, is to pursue improvements in a hundred places at once. However, the insurers seeing results tend to do the opposite. They concentrate on a set of high-impact agentic AI use cases, chosen primarily for their anticipated effect on growth and customer experience, and rewire each domain around these.

In Asia–Pacific, particularly in markets where operations are often more complex to coordinate, a fair question is whether agentic AI can earn its keep. A simple test can clarify whether AI can provide value. For any candidate use case, compare the fully loaded cost of the current human process with the all-in cost of the agentic AI alternative. Then add the additional revenue the agentic use case will likely generate.

Several agentic AI use cases consistently pass the test: those that support growth and conversion, improve claims accuracy, coordinate work across fragmented systems, and build capabilities that can be reused. Insurers in Asia–Pacific markets looking to benefit from agentic AI can prioritize use cases with the greatest growth and experience impact and redesign the following five areas around them:

1. Distribution operations. A typical life insurer in the region runs a large, geographically dispersed advisor force with high turnover; a typical multiline carrier runs several intermediated channels at once, each with its own systems and incentives. Many advisor conversations follow a one-size-fits-all pitch, most product and underwriting questions are escalated to a supervisor, most collateral is assembled by hand, and few sales conversations are reviewed.

By contrast, an advisor copilot could answer product and underwriting questions at the point of sale, personalized collateral and benefit illustrations could be generated in seconds, conversations could be transcribed and assessed, and advisors could receive AI coaching and feedback.

2. Underwriting and pricing. Underwriting requires expert judgment, deep file knowledge, and careful balancing of risk and reward. It is also one of the most expensive, slowest, and least transparent parts of the value chain, relying on manual data collection, subjective, underwriter-dependent decisions, and 12- to 18-month actuarial pricing cycles.

Agentic AI can unlock a different operating model. A team of senior underwriters could orchestrate a fleet of agents that handle intake, data enrichment, guideline checking, draft decisions, and quality review. In turn, the underwriters could focus their attention on judgment calls, relationship management, and complex cases.

3. Policy servicing and administration. For most insurers, high-volume contact centers handle repetitive queries, channels are siloed, endorsement processing is manual, and carriers face low renewal rates with no early-warning signals of customer churn.

This situation could change significantly. Agentic AI could autonomously handle 70 to 80 percent of queries across channels with contextual handoff for exceptions.

In much of Asia–Pacific, applying agentic AI in contact centers is not primarily a cost consideration as many already run at a cost per contact much lower than in other parts of the world. Retention and reach are where the AI advantage lies. AI enables conversational servicing in the customer's own language on the messaging channels they already use.

4. Claims management. Traditionally, claims management is phone-based, documents are reviewed manually, and there can be more than 30 days of turnaround time for P&C claims. The claims process is changing. First notice of loss (FNOL) is shifting from phone to digital, agents can extract and validate information from unstructured claim documents, and fraud detection is improving.

Simple claims already run from FNOL to settlement through automation, while complex claims are likely to stay with a senior adjudicator for final decision-making. Agentic orchestration changes the economics of medium-complexity claims, and its value is concentrated in this middle band of claims. This is where the adjudicator's role shifts furthest, from executing steps to reviewing decisions.

5. Corporate and support functions. Finance, HR, legal, IT, compliance, and procurement account for a large share of an insurer's indirect costs and have been slow to industrialize because the work is document-heavy, judgment-laden, and cycle-bound. This is starting to change. In finance, AI agents can take on much of the transactional and compliance work to free people for strategic planning, forecasting, and enterprise performance management. They can also monitor regulatory feeds and map changes to the processes they affect, a job most insurers still do by hand. And in HR, they can speed up onboarding processes.

To capture these benefits, insurers need to think beyond AI. Technology is rarely the constraint in an AI-led operations transformation. Typically, what holds an AI transformation back is the organization itself. For instance, whether it has aligned the top team on a bold strategy and business outcome, set up the right talent and operating model to execute the transformation, made the right technology and data choices, and committed to the change management required for at-scale adoption. These factors often make the difference between success and unfulfilled outcomes.

Insurance leaders in Asia–Pacific have an immediate opportunity to harness agentic AI to fundamentally transform their operations. Such a transformation is a strategic reinvention of the way insurance work gets done and increasingly involves a choice about where to compete. Carriers able to spread AI infrastructure across a large premium base may compete on scale, while others will need to excel narrowly and partner for the capabilities they do not build.

Asia–Pacific insurers in next-generation growth markets often contend with greater operational complexity than their counterparts in other parts of the world. For example, they operate across multiple languages, legacy systems, distribution channels, product generations, and regulatory regimes, sometimes within a single organization. This complexity makes agentic AI potentially valuable in coordinating people, processes, and technology across the enterprise.

Introduction: From AI pilots to enterprise intelligence

After a decade of investing in digital tools and predictive AI, with measurable but uneven results, insurers may now have a credible path to industrialize the most expensive and judgment-intensive parts of their operations—not just the routine ones. The reason is that predictive, generative, and agentic AI now work well enough together to reconfigure the underwriter's desk, the contact center, the claims journey, and the finance close as connected, intelligent workflows in which AI agents and humans share the work.

In our work with clients, we have seen this shift taking place. Some carriers have moved agentic AI out of pilots and into the core of the business, cutting claims decisions from days to minutes and automatically underwriting most new business. Lower costs are the primary goal, but these carriers also report faster issuance, higher conversion and persistency, and better customer experience.

This report argues that as AI capabilities become widely available, advantage will come less from the technology itself than from how insurers redesign operating models, embed AI in the front line, and build the capability to scale it. The report builds on the findings of McKinsey's AI and insurance research, including the *Global Insurance Report 2025*. It also narrows the aperture of the research, focusing on insurance operations with particular relevance to insurers in Asia–Pacific markets.

When looking at insurance operations through an Asia–Pacific lens, it becomes clear that several regional characteristics affect how agentic AI could be framed, sequenced, and implemented, when compared with a typical Western European or North American AI transformation.

For instance, insurance in next-generation growth markets is expected to show higher premium growth but lower penetration than in more developed markets in the region.[1] Insurance penetration in markets such as India and Thailand remains below the levels achieved in the region's more advanced markets. In 2025, total gross written premium income corresponded to 3.8 percent of GDP in India, compared with 5.1 percent in Thailand, 7.4 percent in Japan, and 7.6 percent in Singapore.[2] Several insurers are working to expand access, improve conversion, and serve the underinsured profitably.

While insurers in more developed markets may be considering AI to reduce costs, those in Asia–Pacific next-generation growth markets could be looking at agentic AI to directly support an expansion agenda through simplified onboarding, multilingual engagement, assisted underwriting, faster issuance, and lower-cost servicing for mass and emerging-affluent customer segments.

In low-cost Asia–Pacific markets, including India and the Philippines, the pure cost-reduction case for AI is not necessarily compelling because many insurers already run relatively low-cost operations where the token cost of applying large language models may not yet be justified for every use case. In these circumstances, agentic AI could be framed as a productivity lever as well as a growth and outcomes lever that can drive faster policy issuance, allow for better cross-selling and upselling, improve persistency, increase loss-cost accuracy, and lead to better claims outcomes. In doing so, the business case shifts from full-time equivalent (FTE) savings to revenue per service interaction, conversion uplift, and combined-ratio improvement.

Asia–Pacific insurers in next-generation growth markets often contend with greater operational complexity than in other parts of the world. For example, they operate across multiple languages, legacy systems, distribution channels, product generations, and regulatory regimes, often within one organization. This complexity amplifies

[1] "India's insurance market set to shine as growth forecast outpaces other major insurance markets," Swiss Re Institute media release, January 19, 2026.

[2] *Allianz global insurance report 2026: The future of insurance in a fragmenting world,* Allianz Research, May 28, 2026.

agentic AI's potential value as an intelligent orchestration layer, seamlessly connecting people, processes, and technology across the enterprise. It can also raise the bar for localization (including vernacular models and onshore data), oversight, and governance.

In parallel, customer expectations in the region are now set by super-apps (mobile or web applications that provide multiple services, including payments and instant messaging) and digital banks, e-commerce, and instant payments. Expectations for the industry have shifted, and insurers are no longer competing against other insurers in terms of experience. That puts pressure on insurers to offer faster FNOL, real-time status updates, proactive servicing, and insurance built into the apps and platforms of other companies. Agentic AI offers a way for insurers to manage the whole customer journey instead of reacting to it one request at a time, and to hold their own in partnerships where the standards for customer experience are set by companies outside insurance.

These factors affect prioritization, use-case sequencing, change management, and the way the AI business case is constructed for Asia–Pacific insurers.

The following observations frame this research:

- **The potential value unlocked from agentic AI in insurance operations in Asia–Pacific is anchored in revenue.** Given the labor-market economics of the region, many carriers already operate at a cost per policy that Western European and North American insurers cannot match. A cost-centric agentic transformation may not yield the expected savings, but agentic AI could move the needle on new business, conversion, persistency, and loss-cost accuracy. That shifts the measure of success from cost per transaction to revenue per customer interaction and combined-ratio improvement, positioning distribution operations as a critical domain for strategic focus.
- **Technology is no longer the bottleneck.** AI can plan tasks, use tools, examine its own outputs, and collaborate with other AI agents and humans across systems. With this level of technological advancement, the barriers to capturing value are now more likely to be organizational than technical. Enterprise vision, talent, operating models, and change management are often needed to reap the benefits of AI.
- **Asia–Pacific is an early adopter with specific needs.** Generally, customer behavior and acceptance of AI shifted earliest in the region, and several Asia–Pacific-headquartered insurers have moved well beyond pilots. Even so, the region has specific characteristics that need to be considered for AI adoption.

Drawing on our work with insurers across Asia, and supplemented by lessons from the Americas, Europe, and the Middle East, this report unpacks the imperative for AI adoption in the insurance sector in Asia–Pacific, specifically in insurance operations (see sidebar "Our research methodology"). It offers a playbook of high-impact agentic use cases that can transform five operating domains: distribution operations, underwriting and pricing, policy servicing and administration, claims, and corporate and support functions. It also provides a framework that can enable organizations to capture value from their AI investments.

Our research methodology

This report draws on McKinsey's work with more than 200 insurers globally, including over 60 in Asia–Pacific, and on the firm's view of more than 300 core processes inside a typical carrier across personal, commercial, health, and life product lines. For example, McKinsey's *Global Insurance Report 2025: The pursuit of growth* provides industry-level data on premium growth and combined ratios across commercial property and casualty (P&C), personal lines, and life in major insurance markets, including the deterioration through 2022–23 and the recovery in 2024. This report synthesizes insights from the *Global Insurance Report 2025*, as well as the *Future of AI for Insurance* series, and our latest engagements with insurance chief experience officers (CXOs) around the world.

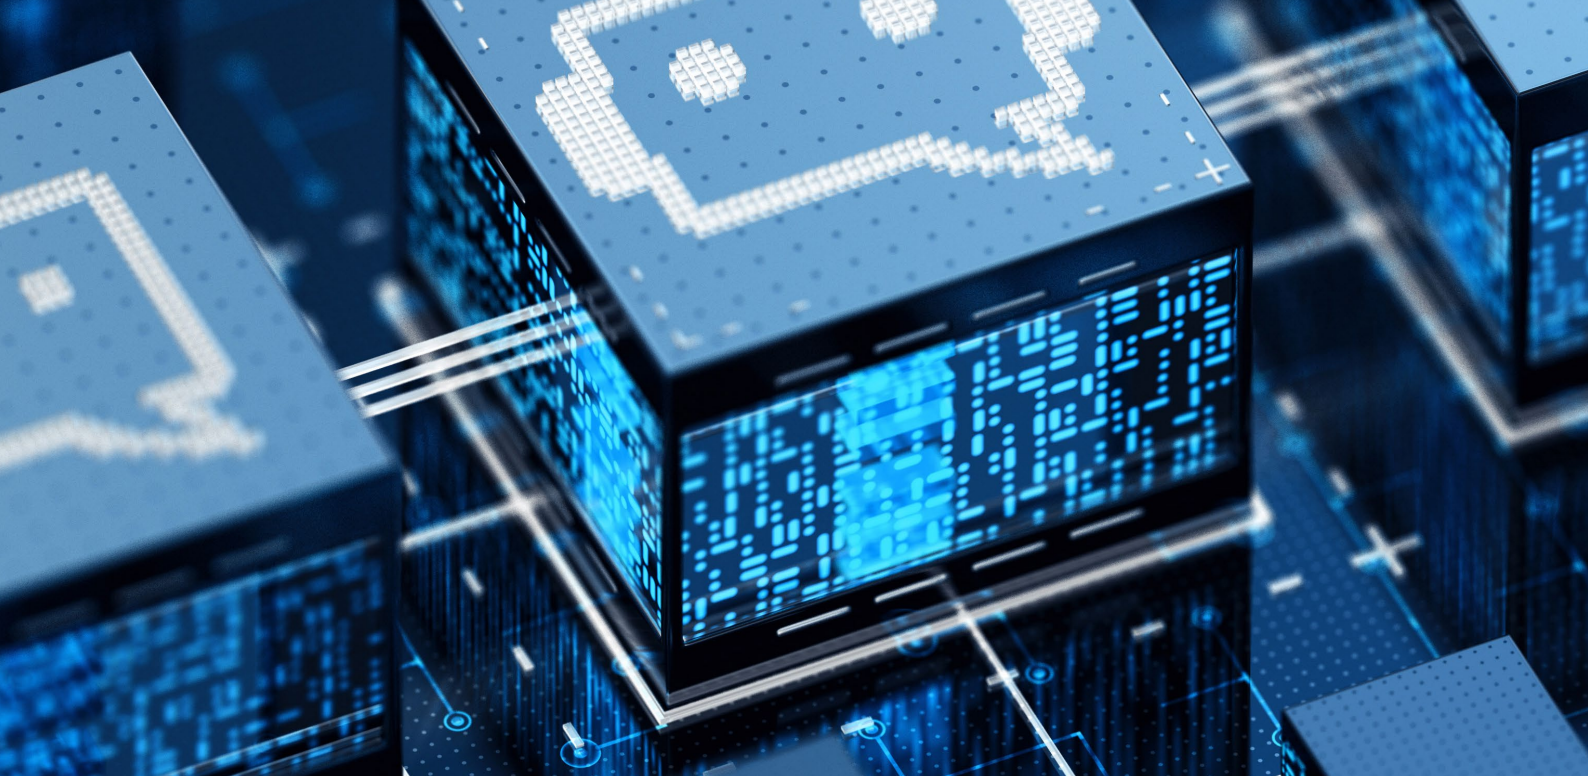

1 The AI imperative for insurance in Asia–Pacific

Globally, insurers face increasing pressure on profitability. In Asia–Pacific, the challenge is particularly consequential. Financial-services penetration has grown rapidly, but insurance protection has not kept pace.

The imperative for insurers in the region is to capture meaningful value from the investments they are making in AI and to use the technology as a catalyst to create competitive distance to peers. Several insurers have embarked on this journey but not at the speed and scale that we think is possible across the region.

Insurers in Asia–Pacific markets looking to reap value from agentic AI can select the use cases with the greatest growth and experience impact and rewire the domain around them. In our experience, the differentiator is AI-first process reimagination. Start with the P&L outcome, redesign the end-to-end journey, define explicit from-to shifts in workflows and human-agent roles, then deploy reusable agents through a common architecture. This helps insurers choose the most suitable AI use cases and move from isolated use cases to end-to-end transformation.

Several forces are now converging to make this moment critical for the sector as carriers assess the value of AI and move from pilot projects to enterprise-wide adoption.

Customer expectations in financial services have reset, especially in Asia–Pacific

Insurance customers in the region no longer benchmark against other insurance companies. They compare their user experience with that offered by super-apps, digital banks, e-commerce, and instant payments. Examples include Alipay and WeChat in China, Paytm and PhonePe in India, PayPay in Japan, Kakao in Korea, and GoTo and Grab across Southeast Asia. These ecosystems, and others like them, have normalized same-second authentication, real-time status updates, and conversational service in local languages. The service implications for an industry where customers generally need to read lengthy documents and wait weeks for claims payments are not insignificant.

With the customer-experience bar now set *outside* the industry, and being reset every quarter, insurance operations need to change. FNOL and servicing have to become real-time, not next-business-day; insurance has to become embeddable inside ecosystem partners, not simply sold through them; and servicing has to become proactive. Agentic AI can deliver these shifts at scale because each requires planning, tool use, and integration across systems.

A 2025 customer-experience survey of more than 2,500 insurance policyholders in India—one of Asia–Pacific's most digitally engaged markets—reveals the scale of dissatisfaction. Key findings include that about 60 percent of survey respondents think premiums are too high; more than 50 percent cite complex policy terms as a pain point; and more than 50 percent considered switching or switched after a poor claims experience (Exhibit 1).[3]

Such severe dissatisfaction is a retention, share-of-wallet, and brand-equity problem. It is also a problem agentic AI is well-suited to address as most of the pain points cluster around servicing friction, opaque communication, and slow, manual claims—the same workflow archetypes where multiagent systems excel.

Exhibit 1

Insurance customers expect simpler, more personalized, and lower-effort journeys.

Customer expectations, % of respondents

Personalized cover that fits needs	**45** say the policy does not suit their needs	**>60** of high-net-worth individuals complain about lack of customization on products
Seamless claims journey	**>45** cite poor claims experience as a pain point	**>50** either considered switching or switched after a poor claims experience
Always-on personalization	**60–65** are discontent with personalization of experience	**>65** of customers are dissatisfied with staff courtesy
Low-effort servicing across channels	**35** are dissatisfied due to poor customer service and lack of support	**~20** cite long and complex paperwork, and lack of digital channels
Cheaper products, clearer value	**55–60** think their premium is too high	**>50** cite complex policy terms as a pain point

Note: Where the source gives a range, the lighter segment shows the upper end.
Source: Insurance Brokers' Association of India (IBAI) Insurance Insights Survey (n = 2,500 retail customers in India), *Leading the path to insurance for all: The broker of the future,* IBAI, July 2025

McKinsey & Company

[3] *Leading the path to insurance for all: The broker of the future,* Insurance Brokers' Association of India (IBAI), July 2025.

The next reset is already visible. Customers in the region have moved from being passive recipients of insurance—found by an agent, sold a product, and served on the strength of a personal relationship—to active choosers who research online and pick channels to suit their wealth and sophistication. The customer journey is shifting from a vertically integrated model, in which a customer deals mainly with a single channel from discovery to claims, to a horizontal digital–physical hybrid in which customers move between channels to optimize for transparency, quality of advice, speed, and convenience.

We have observed that three customer archetypes are emerging: 1) human-touch-oriented customers who start using AI but keep their key interactions with an advisor, at least for complex products; 2) information-driven customers who conduct their initial consultation and planning with AI, then talk to a professional agent, relationship manager, or advisor; and 3) AI-savvy customers who complete the whole journey through digital channels with AI assistance.

Across all three, simple products are increasingly bought directly. The transition to AI-only journeys may take several years, but some segments are already migrating, and insurers will need to understand how customers see the role of each channel—and manage several channels at once.

This makes visibility into AI-generated answers critical. As customers rely on AI for research, what an AI assistant says about an insurer's products increasingly shapes the conversation an advisor or relationship manager has next. A disciplined generative engine optimization (GEO) program typically involves five steps:

- Identify the life moments—starting a family, buying a home, planning retirement—when priority customer segments turn to AI for decisions, and the questions they may ask
- Benchmark the insurer's visibility, positioning, and sentiment against peers across AI engines, markets, and languages
- Strengthen citation-ready owned content and secure a presence in trusted third-party sources such as comparison sites, financial media, and digital partners, in line with regulatory requirements
- Deploy tools to track the visibility and accuracy of AI answers about the insurer's products against approved materials
- Establish a regional GEO team that embeds compliance and provides always-on content updates

Margins are under sustained pressure across most insurance lines

Profit growth has lagged premium growth over the past two decades. Globally, gross written premiums expanded at about 4.9 percent a year since 2005, reaching an estimated $8.3 trillion in 2025, while profit before tax grew at about 4.3 percent a year, to about $580 billion.[4]

Across pockets of Asia–Pacific, the underlying economics of the industry have stagnated. McKinsey analysis shows that between 2018 and 2024, health combined operating ratios failed to improve in half of the region's ten largest markets, P&C combined operating ratios failed to improve in half, and after-tax return on equity in life failed to improve in two-thirds—a sign of stalled profitability rather than cyclical softness (Exhibit 2). The pattern is similar in several large European and North American markets, where loss-cost inflation, claims severity, and tech-debt service have eroded gains from underwriting discipline.

[4] "How AI will reshape the economics of insurance: A CEO's guide to strategy," McKinsey, July 23, 2026.

Exhibit 2

Cost and margin pressure persist in most lines of business across geographies in Asia–Pacific.

Change in performance by line of business, 2018 vs 2024

● Improvement
● No improvement

	Health, combined operating ratios	Property and casualty, combined operating ratios	Life, after-tax ROE
Australia	●	●	●
China	●	●	●
Hong Kong SAR[1]	●	●	●
India	●	●	●
Indonesia	●	●	N/A[2]
Japan	●	●	●
Malaysia	●	●	●
Philippines	●	●	●
Singapore	●	●	●
Thailand	●	●	●

[1]Special Administrative Region of the People's Republic of China.
[2]Data not available.
Source: McKinsey Insurance and Operations practices analysis

McKinsey & Company

This stagnation is a structural problem. Insurance cost ratios are about 10 percent higher globally than they were in 2005. Looking at the Asia–Pacific region overall, cost ratios have declined on average over the same period. But this does not take into account regional variations, as some markets have improved while others have declined.[5]

The global stagnation trend will likely not be resolved by another round of incremental digitization, robotic process automation, or offshoring on top of the legacy stack. It requires a different unit-economics model—one in which a meaningful share of judgment-heavy operational work is executed at the marginal cost of compute, with humans deployed to the cases that genuinely need them. The nuances between Asia–Pacific markets should also be taken into account.

Technology is ready for wider use across the business

For most of the past decade, AI in insurance has largely been used for predictive models such as those for propensity to buy, propensity to lapse, fraud scoring, and micro segmentation. These models worked well with structured inputs, repeatable workflows, and for narrow decisions. Typically, they were not effective when the work was unstructured, judgment heavy, or conversational in nature—which characterizes a great deal of insurance operations work.

Generative AI changed this by making unstructured content easier to manage. A medical report, an attending physician statement, a broker's email, or an inbound voice call could, for the first time, be included and controlled in enterprise systems.

Agentic AI has changed the shape of the work itself. A multiagent system can plan a multistep workflow, use tools and core systems, examine its own intermediate outputs, escalate to a human when needed, and learn and improve from feedback. Simply put, AI can act as a colleague, not merely as an answer engine (Exhibit 3).

[5] "How AI will reshape the economics of insurance: A CEO's guide to strategy," McKinsey, July 23, 2026.

Exhibit 3

Multi-agent systems can execute tasks autonomously within clear guardrails—a step-change from predictive models and single LLMs.

Comparison of predictive models, single LLMs,[1] and multi-agent systems

● Supported
● Not supported

	Predictive models	Single LLM	Multi-agent systems
Examples	• Propensity to lapse • Propensity to buy • Next best action • Microsegmentation	• Summarize a document/content • Ask a question (copilot) • Draft an email or simple content	• Automate complex multistep workflows • Effective two-way conversations • Execute tasks on the user's behalf via access to tools and systems
Capabilities			
Process unstructured data	●	●	●
Interact in human-like language	●	●	●
Reliability	●	●	●
Execute multistep tasks	●	●	●

[1]Large language model.
Source: McKinsey Insurance and Operations practices analysis

The opportunity to gain an early advantage is shrinking

Early movers are creating structural distance from late adopters. Those that stall AI programs may risk competing against insurers that can scale at double or triple their pace, and at roughly half the unit cost. Several Asia–Pacific insurers are already demonstrating what AI scale looks like and are well past the experimental or pilot phase.

Ping An has embedded AI across underwriting, claims, and sales. As a result, life policies are automatically underwritten in seconds. The company notes that 56 percent of life claims are now settled through its Smart Quick Claim channel, which closes a claim in 7.4 minutes on average.[6]

Similarly, the AIA Group states that new business e-submission rates are at 99 percent, with 82 percent of policies being automatically underwritten. Consequently, back-office unit costs decreased by 43 percent from 2020 to 2024.[7]

ZhongAn deeply integrated AI across multiple processes, from underwriting to product design, marketing, claims, and quality control. The company describes its AI middle platform as the engine driving business transformation. The platform deploys about 220 AI agents and has led to efficiency and productivity gains. For instance, in health insurance, the automated case review ratio is above 45 percent, with the fastest cases being settled within 15 seconds. Additionally, about 76 percent of customers receive claim payments within one business day. In auto insurance, AI-based damage assessment can take 116 seconds.[8]

In another example, Zurich Insurance Group has about 460 AI use cases in practical application or in production across the company. For example, in

6 "First quarter financial results for the three months ended March 31, 2025," Ping An media release, April 25, 2025; "Ping An takes lead in Fortune Fintech Innovator Asia list, ranked 1st in insurtech," Ping An media release, November 7, 2024.
7 "Annual results for the year ended 31 December 2024," AIA Group, March 14, 2025.
8 "ZhongAn Online posts robust 2025 results: AI drives quality and efficiency gains, underwriting profit soars over 40%," BigGo Finance, March 20, 2026.

the United States, Zurich uses high-resolution aerial imagery in combination with AI-driven roof scoring to transform property underwriting. The tool helps underwriters assess property risks more accurately and quickly, flagging damage and exposure concerns in seconds. This has enabled better risk selection, faster quoting, and tailored pricing, especially in locations with wind and hail exposure.[9] In the United Kingdom, the group's AI-powered underwriting workbench for commercial insurance has reduced processing times by 78 percent in less than two years—helping underwriters work more efficiently and focus their expertise where it matters most.[10]

These cases demonstrate that there is surging appetite and progress among global insurers for adoption of agentic AI-led transformation of their sales, underwriting, and claims functions. While these developments showcase mature adoption, we have noticed a consistent pattern. Where insurers have committed to enterprise-wide, value-chain transformations rather than a portfolio of isolated experiments, they often generate measurable improvements in productivity, customer experience, and loss-cost outcomes. The followers, meanwhile, may find themselves 18 to 24 months behind.

The discipline of claims, the unstructured intensity of underwriting, and the long-cycle nature of customer relationships in life and health make insurance an unusually high-value use case for agentic AI, provided that organizations move with the decisiveness that the moment warrants.

Early movers are creating structural distance from late adopters. Those that stall AI programs may risk competing against insurers that can scale at double or triple their pace, and at roughly half the unit cost.

[9] *Zurich Insurance Group: Annual report 2025*, Zurich Insurance Group, March 5, 2026.
[10] "Zurich UK cuts commercial underwriting time by 78%," Zurich News & Insights, Zurich website.



2 High-impact agentic AI use cases to rewire domains

Insurance is broad. A typical multiline carrier runs roughly 300 core processes across personal, commercial, health, and life insurance, as well as corporate and support functions—spanning distribution operations, underwriting, policy administration, claims, finance, HR, compliance, and IT. Faced with that surface area, the temptation is to pursue impact in a hundred places at once.

However, the insurers extracting real value tend to concentrate on a set of high-impact agentic AI use cases, chosen primarily for their anticipated effect on growth and customer experience, and sequence them across five operating domains. Individually, each use case is valuable; sequenced across an operating domain, they compound impact by lifting growth (new business, conversion, cross-selling, and persistency) and customer experience as well as efficiency.

The cumulative impact of AI implementation cannot be ignored. McKinsey's estimate puts the economic value generative AI could unlock in insurance globally at between $50 billion and $70 billion, with the highest impact in marketing and sales, customer operations, and software engineering.[11] Domain-level evidence supports this aggregate: insurers running comprehensive domain rewires have reported 10 to 20 percent uplifts in sales conversion, 10 to 15 percent premium growth, 20 to 40 percent reductions in onboarding cost, and 3 to 5 percent accuracy gains in claims.[12]

In Asia–Pacific, particularly in markets where operations can be more complex to orchestrate, a fair question is whether agentic AI can earn its keep. We believe it does—but not everywhere, and not on the basis that most business cases are built on.

We use a simple check. For any candidate use case, we compare the fully loaded cost of the current human process against the all-in cost of the agentic AI alternative (including inference and orchestration, plus the oversight, evaluation, and rework the process will need). We then add the revenue effect—in other words, how much incremental revenue the agentic use case will likely generate.

[11] "The economic potential of generative AI: The next productivity frontier," McKinsey, June 14, 2023.
[12] "The future of AI in the insurance industry," McKinsey, July 15, 2025.

In North America and Western Europe, the cost comparison often carries the case on its own. A shift to agentic AI will often result in cost savings. In much of Asia–Pacific however, the cost comparison does not make the case for agentic AI, given much of the region's relatively lower labor costs. But the revenue growth that it could generate often tips the scales. Essentially, that difference could determine an insurer's investment in specific agentic AI use cases.

Several use cases consistently clear the bar when applying the screen. First, growth and conversion because the AI upside is measured in new business, persistency, and claims adjudication. Such use cases span advisor enablement, simplified and multilingual onboarding, assisted underwriting, and faster issuance. Second, accuracy-led use cases in claims are often worth investment, as a reduction of a few percentage points in loss-cost or fraud leakage outweighs the investment in reducing servicing cost. Third, agentic AI can provide orchestration across fragmented estates, reducing the need for human coordination across teams that demand complex functional integration and coordination. The final set of use cases that typically pass the test involves reusable foundations (including document ingestion, a semantic layer, and the agentic backbone). These may look expensive against a single market but are inexpensive when amortized across a group.

There are also instances where agentic AI may not be the answer. High-volume, low-complexity transactional work that is already performed at low cost is the clearest example. Once oversight and rework are counted, the agentic unit cost can land at or above what the insurer pays today, and the honest answer may be to wait for inference costs to fall further.

Nor is there much return in applying language models where a deterministic rule or an existing predictive model already performs well. Agentic approaches earn their keep on unstructured, judgment-laden work—using them elsewhere may add cost and review burden without adding accuracy.

Per-market point solutions are another common short-term option that may not yield long-term benefits. A use case built for one market with no path to reuse it across the group rarely repays its build cost in a midsized Asia–Pacific market. Core-system replacement should not be a precondition for using agentic AI, as the insurers moving fastest are layering agents over the estate they already have.

There may not be much value in blindly including AI in insurance operations. Value is more likely to come from redesigning insurance operations around what AI and agents now make possible.

The practical consequence of these factors is that insurers in Asia–Pacific are faced with a different sequencing decision from the one a Western carrier might make. The five domains that follow are ordered accordingly: distribution operations first, because that is where the growth case is strongest and the Asia–Pacific operating gap widest, and corporate and support functions last, where the value is real but rests on complexity rather than on scale (Exhibit 4).

There may not be much value in blindly including AI in insurance operations. Value is more likely to come from redesigning insurance operations around what AI and agents now make possible.

Exhibit 4

Priority agentic AI use cases are spread across five operating domains; the Asia–Pacific context influences which domains matter most.

Priority agentic AI use cases by operating domain

Domain	Priority agentic AI use cases	Where the value comes from	Why the domain matters more in Asia–Pacific
Distribution operations	• Microsegmented campaigns • Advisor copilot • Conversation transcription • AI coaching	Growth: Reach, conversion, cross-sell, and advisor retention	The workforce is typically very large and dispersed, and there is high churn, especially in the tied agency model
Underwriting and pricing	• Automated ingestion of unstructured submissions • AI risk scoring with explainable outputs • Guideline and rule checking • Draft decisions with human review • Frequent, governed pricing reviews using approved models and claims data	Effectiveness: Accuracy and consistency Efficiency: Capacity and turnaround	AI could be beneficial for risk discovery and expansion, especially given shorter case histories, fragmented systems, and inconsistent data
Policy servicing and administration	• Conversational AI agent in voice and messaging channels • Always-on agentic assistant copilot • Automated endorsement and document processing • Proactive lapse and renewal intervention	Growth: Persistency and retention Efficiency: Cost to serve	The bar is raised in customer experience expectations, insurers compete against super-apps that offer service in the vernacular
Claims management	• Digital and voice-guided FNOL[1] • Intelligent document extraction and validation • Fraud detection • Autonomous adjudication of medium-complexity claims, with a human review and escalation	Effectiveness: Loss-cost accuracy and fraud leakage Efficiency: Cycle time and handling cost	The argument for agentic AI is claims accuracy, rather than claims automation
Corporate and support functions	• Automated reconciliation, close, and management reporting • Regulatory-change monitoring • Knowledge agent for onboarding and policy queries	Effectiveness: Control and reporting quality Efficiency: Close cycle and transactional effort	Insurers are typically multientity, multicurrency, multijurisdiction

[1]First notice of loss.
Source: McKinsey Insurance and Operations practices analysis

McKinsey & Company

Domain 1: Distribution operations

A typical life insurer in the region runs a very large, geographically dispersed, and high-churn advisor force; a typical multiline carrier runs several intermediated channels at once, each with its own systems and incentives. Leads tend to be pushed on volume rather than propensity; many advisor conversations default to a generic one-size-fits-all pitch; most product and underwriting questions escalate to a supervisor, most collateral is assembled by hand, and few sales conversations are ever reviewed.

None of that is a product problem or a pricing problem. It is a distribution-operations problem, and it is the one agentic AI is best suited to solve—which is why distribution belongs in an operations playbook rather than being treated as a separate commercial agenda.

AI has the potential to turn the current situation into one where an advisor copilot answers product and underwriting questions at the point of sale, personalized collateral and benefit illustrations are generated in seconds, most conversations are transcribed and graded, and the advisor receives AI coaching and feedback.

The competitive stakes are rising fast in distribution operations as insurance remains overwhelmingly intermediated. For instance, in the United States, about 85 percent of P&C premiums and 95 percent of life premiums still reach the customer through agents, brokers, and managing general agents, and nearly half of North American customers already use AI somewhere in their personal insurance-buying journeys.[13] As customers' AI assistants begin to monitor renewals, compare cover, and recommend switches, differentiation becomes increasingly driven by the quality of the recommendation and the data powering it. We anticipate that disintermediation will move fastest in commoditized personal lines. In midmarket commercial and specialty risk, AI is more likely to compress cost and lift advisor productivity than to replace the advisor, a distinction that matters in Asia–Pacific, where agency and bancassurance carry the bulk of life distribution.

Agentic AI could potentially transform an advisor's day across the customer servicing life cycle (see sidebar "Frontline sales enablement at a leading Indian life insurer"). Solutions such as hyperpersonalized targeting engines, advisory copilots, and sales coaching assistants could boost an advisor's productivity in identifying the right customers, holding conversations with them, and ultimately converting the sale and building the long-term relationship (Exhibit 5). Early evidence is encouraging. AI transformations are already delivering 10 to 20 percent higher agent productivity.[14]

Exhibit 5

Agentic use cases span the advisor's day—from lead to conversation, offer, and coaching—rather than the moment of sale.

Three capabilities layered across the frontline sales journey, each compounding the others

Target customers	Conduct research	Prepare for the call	Hold the conversation	Prepare the offer	Track performance	Act on feedback
Relevant leads	Questions answered from one knowledge base	Segment-specific sales script	Personalized script guides the conversation	Tailored collateral and benefit illustrations	Daily grades and start-of-day briefings	AI recommendations plus team-leader input

① Hyper-personalized campaigns

Microsegments built on client profile, loyalty status, and risk exposure drive targeted campaigns and next-best-action leads—the growth engine of the rewire.

② Knowledge base and advisor copilot

Draws on a repository of over 1,000 internal documents, answers product and underwriting questions at the point of sale, and generates personalized collateral in seconds.

③ Sales coaching assistant

Voice-based review grades every conversation rather than a sample, then returns tailored coaching alongside team-leader feedback.

Capabilities overlap by design: each one is more valuable because the others are in place.

Source: McKinsey Insurance and Operations practices analysis

McKinsey & Company

[13] "How AI will reshape the economics of insurance: A CEO's guide to strategy," McKinsey, July 23, 2026.
[14] "How AI will reshape the economics of insurance: A CEO's guide to strategy," McKinsey, July 23, 2026.

Frontline sales enablement at a leading Indian life insurer

An Indian life insurer set out to lift advisor productivity across a frontline base of about 300,000 users. The diagnosis was operational rather than commercial. A fragmented application landscape gave advisors a limited and incomplete view of the customer journey. The application redirected product and underwriting queries to the advisors' supervisors, which led to generic customer conversations.

The rewire began with a blueprint phase where field research and focus groups with advisors were used to prioritize use cases and validate a working prototype before a delivery partner was selected. The two highest-value capabilities—product query resolution and on-demand personalized collateral—were piloted to a controlled user group. The pilots were tested with continuous feedback loops before scaling up a sales-pitch assistant that managed objection handling and query resolution, provided talking points for meetings, and offered proactive nudges to advisors.

Other insurers in the region can adapt this design choice for their needs. The rewire highlights a deliberate effort to codify user stories and reusable knowledge bases so the same capabilities can later be extended into underwriting, servicing, and claims.

Domain 2: Underwriting and pricing

Underwriting has historically been the prestige discipline of insurance as it requires expert judgment, deep file knowledge, and careful balancing of risk and reward. It has also been one of the most expensive, slowest, and least transparent parts of the value chain. Typically, it relies on manual data collection; subjective, underwriter-dependent decisions; and 12- to 18-month actuarial pricing cycles. Advances in AI have materially expanded what is possible across complex underwriting workflows (see sidebar "A large commercial property carrier improves underwriting"). Yet much of today's underwriting process looks much as it did a decade ago. Queues of broker submissions and underwriters reading emails, extracting facts from documents, chasing missing information, and manually moving cases between various stages over days, sometimes weeks.[15]

Agentic AI can unlock a different operating model based on automated multisource data ingestion; AI risk-scoring with explainable outputs and human review based on risk and case complexity; and more frequent, governed pricing reviews using approved models and claims data. Underwriters can rely on agentic AI for specific tasks across processes from submission to prioritization, risk assessment, quoting, issuance, and portfolio monitoring (Exhibit 6).

This new way of working can free up time for the underwriters to focus their attention on judgment calls, relationship management, and complex cases (Exhibit 7).

[15] "The future underwriting operating system: From inbox to AI nerve center," McKinsey, June 8, 2026.

Exhibit 6

In underwriting, AI can unlock efficiencies and improve the quality of decision making.

Illustrative commercial property example

● Underwriter actions ● Agent/AI/tech enablement

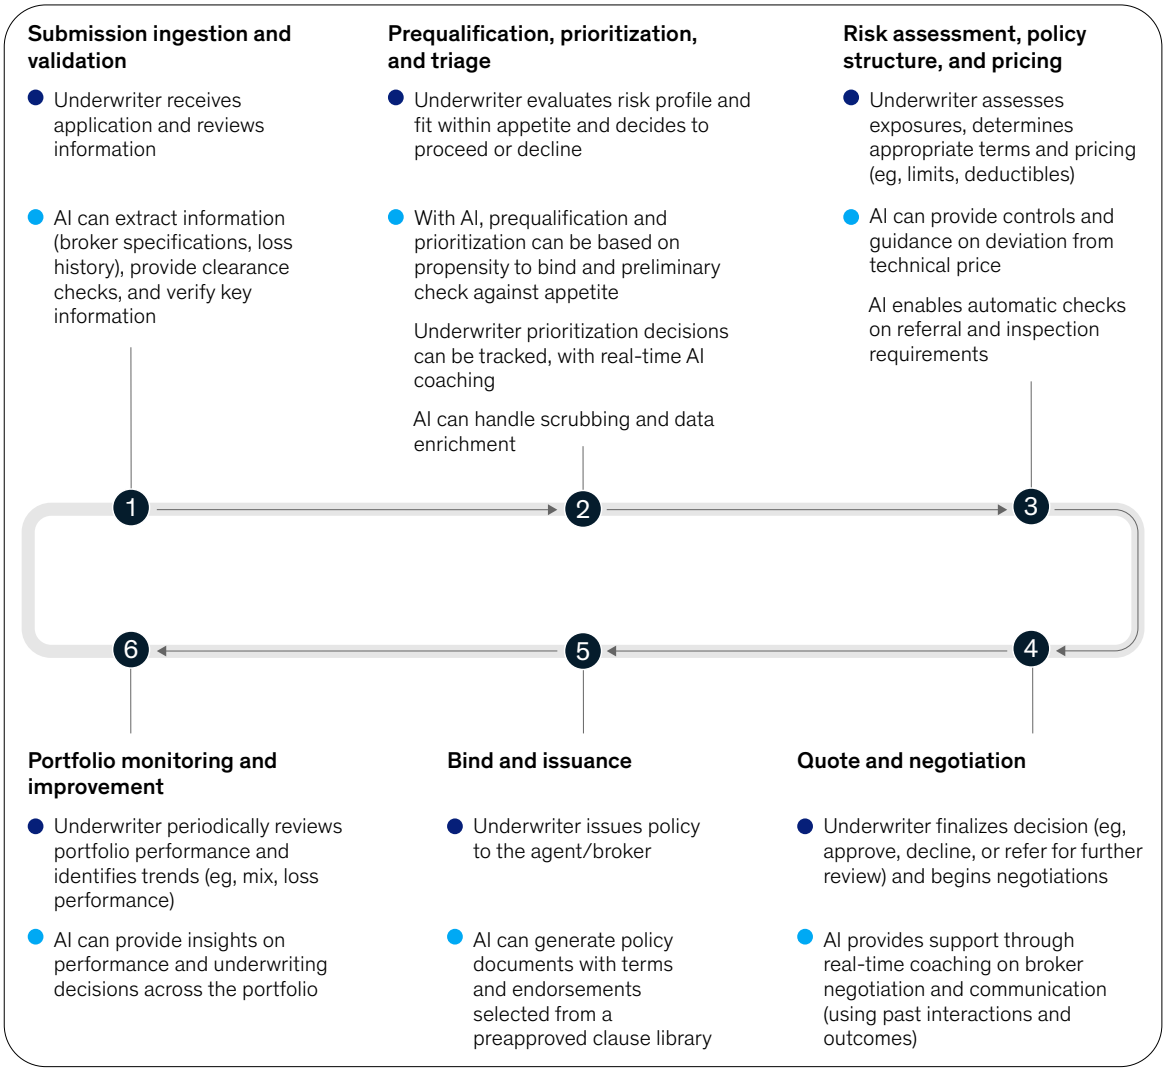

Source: McKinsey Insurance and Operations practices analysis

McKinsey & Company

Exhibit 7

An AI-enabled operating model can free up underwriting time.

Time spent on selected work activities by insurance underwriters, midpoint scenario with gen AI acceleration, % of time

	Current operating model	AI-enabled operating model	Current operating model / AI-enabled operating model
Information collection and verification	Manually reviewing submission information across disparate documents	Leveraging AI-enabled ingestion, structuring, and enrichment of submission information	35 / 20
Risk assessment and pricing	Using a judgment-based approach to exposure assessment, pricing, and policy structuring Receiving unclear feedback	Leveraging AI-enabled initial assessment of risk and automated quality-related flags	45 / 35
Quote, negotiation, and bind	Spending time on quote preparation and iterations during negotiation Applying ad hoc, nonstandard approaches to negotiation	Streamlining the quote and bind process with AI assistance and minimal instructions Gaining AI-enabled coaching support on broker negotiation	20 / 15

Note: Activity mapping based on McKinsey Global Institute (MGI) taxonomy.
Source: MGI Automation Model June 2023, selective HR data input on roles, activities, time spent

McKinsey & Company

In Asia–Pacific, the volume of insurance sits in retail, life, and health rather than in large commercial risk, which changes what to build first. The constraints are medical evidence in multiple languages and scripts, attending physician statements and laboratory reports that arrive as scans rather than as data, and underwriting manuals that differ by market and by product generation. This makes document ingestion and evidence summarization the highest-return starting point, ahead of the submission-triage use cases that dominate commercial P&C programs in London and New York and in the commercial hubs of Australia, Japan, and Singapore.

Partner data is a second underwriting lever. For mass-market and emerging-affluent customers reached through banks and digital platforms, insurers can use consented customer profiles and look-alike models to preapprove standardized products with simple eligibility and simplified underwriting—avoiding medical evidence altogether for a large share of cases and enabling straight-through issuance inside the partner's app.

A large commercial property carrier improves underwriting

A large commercial property and casualty (P&C) insurer set out to expand into the middle market while increasing its speed to quote and underwriting consistency. The previous process was highly manual for tasks such as clearance, statement of values (SOV) cleansing, and form review and was fragmented across more than 30 underwriting tools.

The redesign introduced automated intake and data-driven triage, differentiated quote flows, embedded quality checks, AI-assisted form review, and a unified underwriter workstation.

Reported impact includes an uplift in underwriting productivity, less than 24-hour quote turnaround for middle-market submissions, and a 60 percent reduction in underwriting touch time. The insurer is now on a path to deploy AI across all lines of business.

Domain 3: Policy servicing and administration

For most insurers, contact centers remain attrition-heavy and reactive. Customers get in touch with an endorsement request, a premium query, a renewal question, or a complaint, and the cost-to-serve scales linearly with volume. Typically, high-volume contact centers handle repetitive queries, channels are siloed, endorsement processing is manual, and carriers face low renewal rates with no early-warning signals of non-renewal or customer churn.

Agentic AI could change this situation significantly, autonomously handling 70 to 80 percent of queries across channels with contextual handoff for exceptions (see sidebar “An AI voice agent at a motor insurer”).

Agentic AI is changing the cost curve in three ways simultaneously. Routine interactions are absorbed by an AI agent before they reach a human, the human agents that remain are supported by an always-on copilot, and the entire complaint pipeline becomes a feedback loop that proactively addresses root causes (Exhibit 8).

In Asia–Pacific, applying agentic AI in contact centers is not primarily a cost consideration. Contact centers across much of the region already run at a much lower cost per contact than in other parts of the world, so a business case built on average handling time tends to disappoint. Retention and reach are where the AI advantage lies, enabling conversational servicing in the customer’s own language on the messaging channels they already use. AI can ensure an intervention before a lapse, moving the target metric from cost per contact to persistency and revenue per interaction. It also raises the bar on language coverage because a servicing agent that performs well in English and poorly in other languages is unlikely to be trusted with customers at scale.

Exhibit 8

An AI-enabled servicing model spans proactive resolution, automated interactions, and agent assist—with regulatory and feedback loops baked in.

AI agents across policy servicing operations

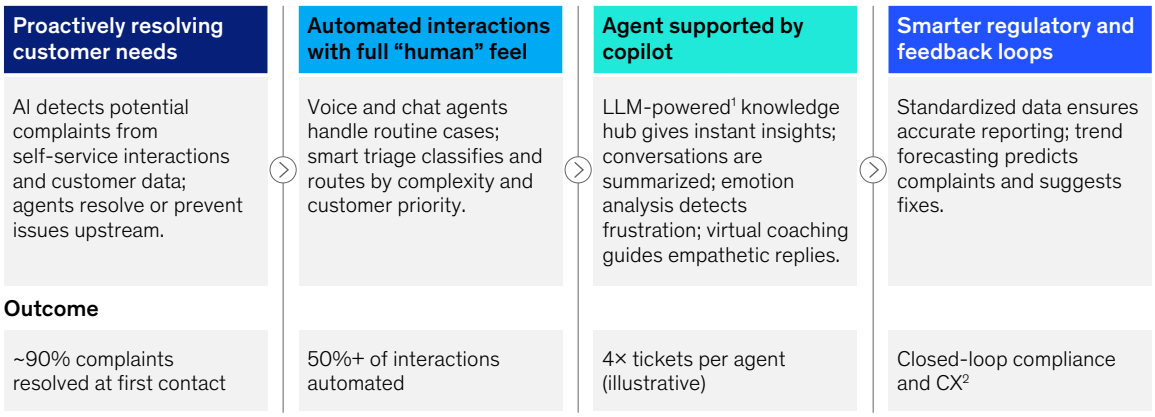

[1]Large language model.
[2]Customer experience.
Source: McKinsey Insurance and Operations practices analysis

McKinsey & Company

An AI voice agent at a motor insurer

A well-established motor insurer in a developed market had several issues that it wished to solve: an inefficient interactive voice response (IVR) system, manual triage that caused a bottleneck, and multiple touches and cold transfers that led to repeated identification and verification.

Within weeks, and inside its existing technology environment, the insurer stood up a proof-of-concept AI voice agent that asked for and understood the reason for the call, differentiated existing claims from new claims, looked up policy data, accepted in-conversation corrections, and routed the call to the correct team.

This capability is now enabling further versions of the technology to be developed to create initial customer engagement at first notice of loss (FNOL), warm transfers to experts within claims, and some agentic handling across the workflow.

Domain 4: Claims management

Traditionally, claims management is phone-based. Claims documents are reviewed manually, and manual adjudication means that turnaround time for P&C claims can be more than 30 days.

The claims process is being disrupted to an extent that was hard to imagine a few years ago as FNOL moves from phone to digital and "always-on" capabilities become the norm (Exhibit 9).

Exhibit 9

In claims, a hybrid agent and adjuster workforce can achieve 'always-on' capabilities.

Illustrative and simplified automotive bodily injury example

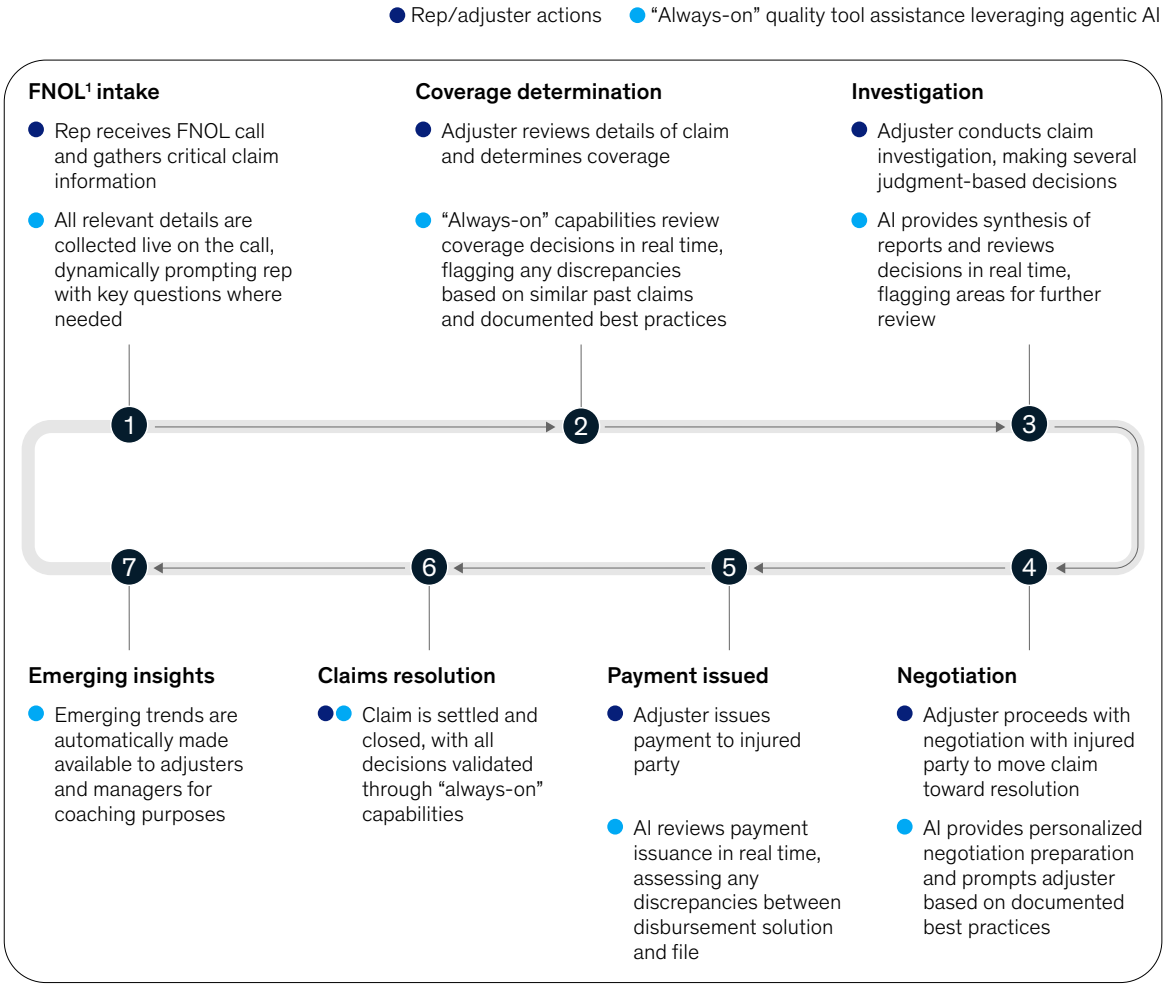

[1]First notice of loss.
Source: McKinsey Insurance and Operations practices analysis

McKinsey & Company

For insurers in Asia–Pacific, the most significant benefit is accuracy. Health portfolios in several markets carry material leakage through inconsistent medical-necessity assessment, duplicate billing, and provider-side inflation. These issues involve document-intensive, judgment-laden decisions to which agentic systems are well suited. And because a percentage point of loss ratio recovered on a health book is typically worth more than the entire cost of processing those claims, claims accuracy rather than claims automation is the primary argument for agentic AI in the claims management domain.

The complexity of a claim determines how much human involvement is required—and where agentic AI can play a role. Simple claims already run from FNOL to settlement through existing automation and are often processed in a straight-through manner. At the other end of the spectrum, complex claims are likely to stay with a senior adjudicator for final decision-making. Between these two scenarios is the medium-complexity claim band, where agentic AI's value is concentrated. Here, subprocesses such as triage, liability, remedy, and repair could run as agentic flows while the human in the loop reviews the output. That is where agentic orchestration changes the economics, and where the adjudicator's role shifts furthest—from executing steps to reviewing decisions.

As in other domains, an AI-enabled workforce can save time, freeing up insurance claims adjusters to focus on other activities (see sidebar "Transforming claims at an East Asian P&C insurer").

Transforming claims at an East Asian P&C insurer

The claims organization of a leading P&C insurer in East Asia was stretched close to breaking point. Handlers were consumed by day-to-day caseloads, strategies and development plans were drawn up separately by each claims center, and every new capability ran into the same wall: changes to a large, monolithic core system. Rather than settle for incremental fixes, the insurer committed to rebuilding its claims operations from the ground up over three years.

The ambition was bold: to free up roughly a quarter of claims capacity and deliver payment-accuracy improvements worth well over $100 million a year. The transformation started with the organization. Business and digital staff were brought together in cross-functional pods, product owners were given real decision rights, and initiatives moved forward in weekly cycles. Every change was first tested in a "model claims unit," refined based on what was learned, and only then rolled out nationwide. Technology choices changed too. By building its new FNOL platform in-house and outside the core system, the insurer cut development costs to about a quarter of the original estimate.

The results are evident. Initiatives that are already live—automated handler assignment, AI-generated case summaries, digital FNOL, enhanced fraud detection, and new rules for repair-shop referrals and payment calculation—have generated $30 million a year in payment-accuracy improvements.

Domain 5: Corporate and support functions

Finance, HR, legal, IT, compliance, and procurement collectively absorb a large share of an insurer's indirect cost and are historically among the slowest to industrialize because the work is document-heavy, judgment-laden, and cycle-bound. Agentic AI changes this situation. Specifically, in finance, transactional and compliance work—about 60 percent of the function's effort—can be automated, redirecting capacity toward strategic planning, forecasting, and enterprise performance management (see sidebar "A diversified insurer builds an AI-powered finance function"). In our experience, finance teams are delivering the most value with AI in strategic planning and control, cash and working-capital management, and cost optimization.[16]

Agentic AI can help insurers move from a seven- to 15-day month-end close to automated reconciliation and reporting with a three-day close. It can also replace manual reporting and tracking of regulatory changes with a system in which AI monitors regulatory feeds and auto-maps changes to the affected processes. Onboarding in corporate and support functions can also be slow. Now an AI knowledge agent can reduce employee ramp-up time by 40 to 60 percent.

Group structures make these advances more valuable in Asia–Pacific than the cost base alone suggests. A regional insurer closing its books across eight or ten legal entities, in as many currencies and reporting regimes, spends a disproportionate share of finance effort on consolidation and reconciliation rather than on analysis. The same multiplicity applies to policy administration and management reporting. Where a single-market carrier sees a modest efficiency case for AI, a multimarket group could use agentic AI to solve a coordination problem that grows with every market it enters.

A diversified insurer builds an AI-powered finance function

A leading insurance firm used agentic AI to rewire its financial planning and analysis processes and strengthen its decision-making. Before the rewire, its finance processes were highly manual and focused on backward-looking activities. For example, about 70 percent of the function's capacity was spent pulling, cleaning, and analyzing data in a manual reporting environment. Forecasts were built on spreadsheets and email inputs, and while the firm had plenty of dashboards, these were not linked or integrated with other functions.

The rewire followed three steps to reimagine financial planning and analysis end-to-end instead of tweaking technology around the edges or bolting AI onto broken processes: define the AI-first future state and design processes accordingly; design the core capabilities needed to enable the future state, including a new operating model and data and tech architecture; and develop a 24-month delivery roadmap and change management plan.

Ultimately, finance became a strategic advisor to the business, delivering proactive, data-driven insights, powered by AI and modern technology.

Impact included a reduction in the close-to-insight cycle from 16 business days to one business day, a 40 percent reduction in planning cycle time, a 50 percent increase in planning and analysis capacity, and a 50 percent improvement in forecast accuracy.

[16] "How finance teams are putting AI to work today," McKinsey, November 3, 2025.

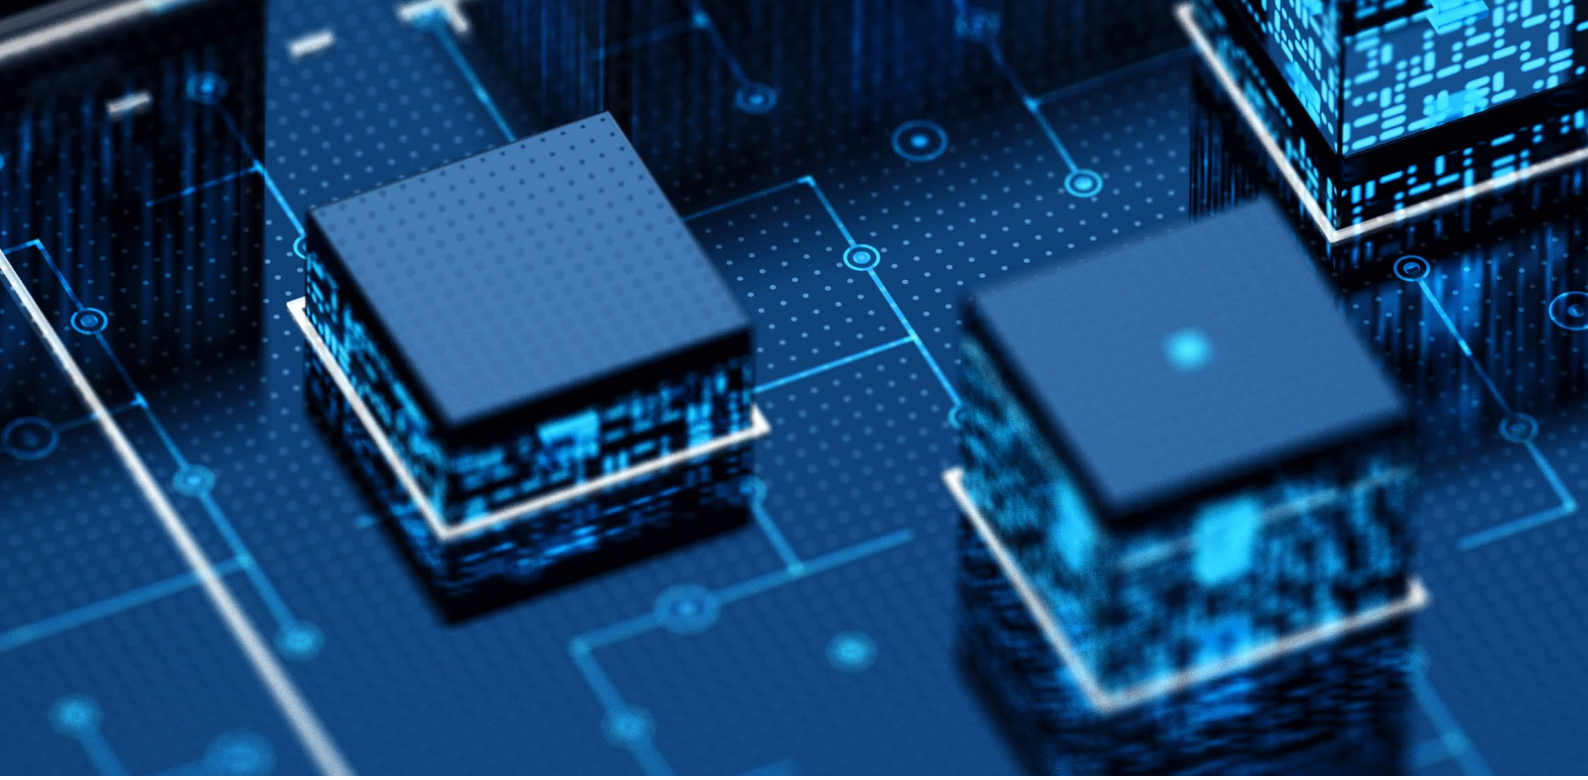

3 The six dimensions of the Rewired framework for at-scale impact

We've observed that technology is rarely the constraint in an AI-led operations transformation. Typically, what holds an organization's AI transformation back is the organization itself. In our work with insurers, the six dimensions of the McKinsey Rewired framework consistently separate the carriers that capture value at scale from those that spend without it.[17] This chapter looks at each of the dimensions in turn: strategy, talent, operating model, technology, data, and adoption and scaling (Exhibit 10).

Technology is rarely the constraint in an AI-led operations transformation. Typically, what holds an organization's AI transformation back is the organization itself.

[17] Eric Lamarre, Kate Smaje, Robert Levin, Alex Singla, and Alexander Sukharevsky, *Rewired: McKinsey's playbook on how leading companies win with technology and AI*, Wiley, Apr 2026.

Exhibit 10

Six dimensions of the Rewired framework are critical for a successful AI-led operations transformation.

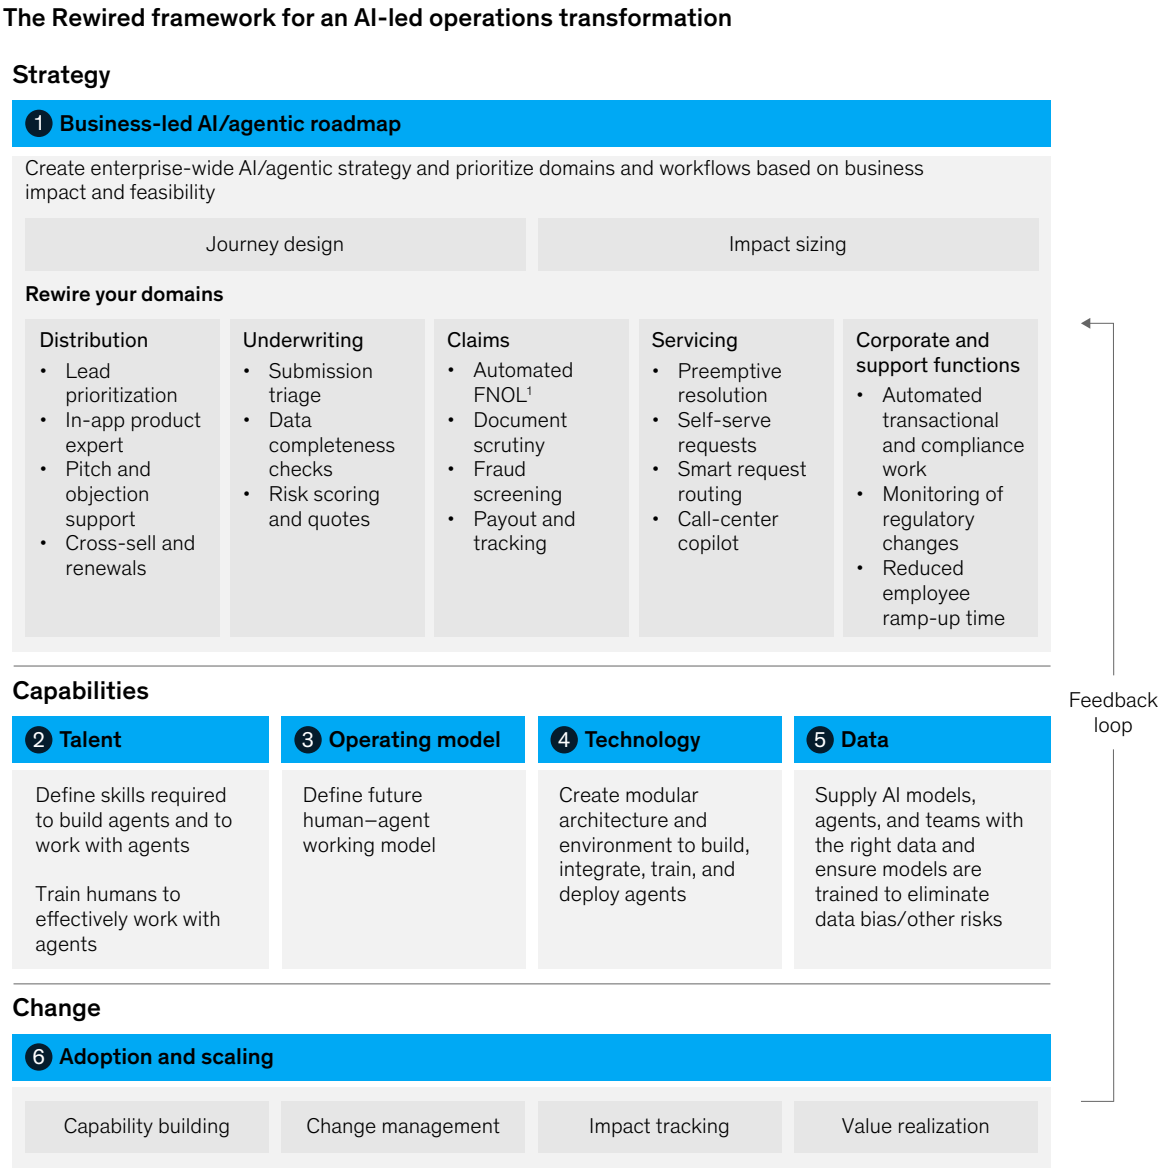

[1]First notice of loss.
Source: Eric Lamarre, Kate Smaje, Robert Levin, Alex Singla, and Alexander Sukharevsky, *Rewired: McKinsey's playbook on how leading companies win with technology and AI*, Wiley, Apr 2026

McKinsey & Company

1. Strategy

Most successful transformations begin with a vision anchored in business value. Ideally, the vision should set explicit outcome targets, prioritize three to five high-value domains for end-to-end transformation, and define owners and milestones with the same rigor as a major M&A integration. And the mechanics of putting the vision into action matter. A one-page CEO and COO ambition statement envisioning what the insurer will be famous for in 24 months, in language a frontline branch manager would recognize, can make the vision accessible. Without it, AI investment tends to scatter across functions and value compounds slowly.

For Asia–Pacific insurers, the ambition statement could also name the measure the business case will be judged on. Where the ambition is growth, persistency, and loss-cost accuracy—rather than productivity unlock—use cases should be chosen accordingly; otherwise, the insurer may risk not seeing the desired returns on agentic investment.

2. Talent and organization

The future agentic insurance team will likely be a mixed-mode team, where three workflow archetypes coexist inside the same operations organization: 1) human-led, agent-enabled; 2) agent-led, human-in-the-loop; and 3) fully agentic within defined limits, with named human accountability and escalation paths. Today, the leadership task is to map the future-state workflows, identify the critical skills required to deliver them, assess talent gaps and surpluses, and develop a workforce reconfiguration strategy—coupled with change management and skilling. This is, in our experience, the most consistently underresourced workstream of an AI-led transformation.

As routine execution moves to agents, demand rises for technology and analytical talent. The tools now available to engineers, such as those for code generation, agentic development environments, and reusable components, shift the premium toward experts who can orchestrate agents while pairing deep domain knowledge with AI fluency. At the same time, demand falls for parts of claims, sales support, and policy administration. The implication is not headcount reduction but reskilling and redeployment at scale. We expect that most roles will be reconfigured rather than removed.

In Asia–Pacific, the workforce constraint is sharper than the global picture suggests. The engineering and AI talent an insurer needs is concentrated in a handful of city markets and is being bid for by technology firms, banks, and global capability centers. People who can bridge insurance domain knowledge with AI fluency are scarcer still. Most insurers we work with resolve this with a hub-and-spoke model: a regional center that builds and governs reusable capability and small translator teams in each market accountable for adoption, language, and local process fit rather than for construction. Where every market staffs its own full team, we have observed that the talent can thin out and the AI build slows down.

3. Operating model and governance

Transformation execution can be orchestrated through cross-functional AI "garages"—agile teams of domain experts, data scientists, engineers, change-management specialists, and transformation leaders, with explicit mandates to test, iterate, and scale AI across the enterprise. We have seen this model unlock speed (the first minimum viable product is often live in six to eight weeks) and quality (because the garage is accountable end-to-end for adoption, not just delivery).

We have found that several governance principles consistently work well: single accountable owners by domain; joint CEO–COO sponsorship; an AI value-capture office that monitors the transformation portfolio with the same discipline as the office of the CFO; and a clear policy on build-versus-buy that reserves proprietary development for genuinely differentiating capabilities.

In practice, this often takes the form of a business-led, centrally steered model, supported by a center of excellence and executed by domain-based, cross-functional AI squads—paired with an explicit goal of role-relevant AI fluency for every employee.

4. Technology

An AI-led transformation does not require an insurer to start the technology stack from scratch. Most carriers have meaningful existing investments to leverage. What is required is the thoughtful introduction of modular, reusable components—agents, orchestration layers, prompt and policy libraries, data pipelines,

and observability—that allow rapid deployment and integration without overhauling legacy infrastructure.

As they scale, the most effective insurers put in place a composable agentic architecture made up of modular agents built on shared services, orchestrated through well-defined patterns, and connected to core systems (policy, claims, and customer relationship management [CRM]) through a governed gateway and tool registry. Underneath sits what we call the agentic backbone, the platform layer that lets agents run safely at scale. The backbone enables agent runtime, ensures observability and monitoring, and provides guardrails for access management. This backbone is what separates AI agents that scale and generate value from pilot projects. Once the agentic backbone is in place, the marginal cost and time-to-impact of each new use case fall sharply.

Our scaling principle is simple: Reimagine once, build common capabilities once, then configure and reuse—as opposed to rebuilding for every market or line of business.

For a multimarket Asia–Pacific insurance group, investing in the agentic backbone is a significant decision. The same use case, rebuilt eight times against eight policy administration systems in eight different languages, is likely to exhaust budget before scale is achieved; however, with an agentic backbone, the use case can be built against a common orchestration layer with market-specific configuration, which improves the unit economics and makes it cheaper with every deployment.

5. Data

Data is the second half of the technology equation and increasingly the differentiator between transformations that deliver value and those that do not. Agentic systems are only as good as the context they can draw on. Accordingly, leading insurers are investing in data products and, critically, a semantic layer and knowledge graphs that give agents a consistent, governed view of business meaning drawn from policies, coverages, customers, and claims. This grounding materially improves accuracy and reduces hallucination, and, because one semantic layer serves many agents and workflows, it compounds in value as more use cases are built.

Two regional realities make getting data right harder and more valuable. Agents have to reason over evidence in multiple languages and scripts, much of it arriving as scanned images rather than as data, so extraction quality varies by market in ways a single-market carrier never has to manage. And product estates are deeper. Asian insurers often have many generations of products still in force, each built on different systems and rules. As a result, the same term—such as premium, benefit, or policy value—can mean different things across products. A semantic layer that reconciles those definitions once is worth more in these circumstances precisely because there is more to reconcile.

6. Adoption and scaling with new ways of working

Adoption does not happen by default. Any AI transformation is more than a tech transformation, so carriers need to determine the level of change management needed to ensure a smoother transition to AI use. A simple four-part approach has served our clients well:

Role modeling. Essentially, seeing is believing. Leaders can use AI visibly in their own work, champion early successes, and support peers in their AI use.

Communicating. Inspiration is key. The organization can clearly state its guiding principles for AI along with a bold aspiration and compelling employee testimonials.

Capability building. This entails expanding skills as well as mindsets. Insurers could consider investing in personalized learning for employees via digital platforms as well as cohort-based skills-building programs.

Reinforcing and sustaining momentum. This is where change gets hardwired. Organizations can provide just-in-time support for employees' AI queries across all workflow channels, offer an AI rewards and recognition program, and track adoption as well as root causes for lack of adoption.

We have observed a far higher likelihood of success if all four elements are used together to drive change.

How to get started

Insurance leaders in Asia–Pacific have an immediate opportunity to harness agentic AI to fundamentally transform their operations. The upside is concrete: breakthrough gains in efficiency and customer experience, improved loss-cost accuracy, and an enduring competitive advantage in a market where the bar for digital experience is set outside the industry. The gap between movers and laggards is already showing up in market value. Across the industry, AI leaders have delivered about six times the total shareholder returns of laggards.[18]

As a first step, we suggest that leaders shift their mental model from a technology-first to a business-first outlook. A bold, business-led vision anchored in tangible outcomes—customer experience, risk reduction, or cost efficiency—is the North Star that guides choices across the five operating domains. Once domains are prioritized, attention can turn to the specific processes inside each that cause the most friction or value erosion, for example manual submission processing in underwriting, IVR-based FNOL in claims, manual reconciliations in finance, and regulatory tracking in compliance.

Many processes share the same building blocks, such as document understanding, case routing, decision orchestration, and conversational interaction, so the early use cases in a domain make the next ones faster to deliver. Sequencing the work this way maximizes return on investment and lets governance and data discipline mature along the journey, building organizational confidence in the transformation.

Ultimately, an AI-enabled transformation is not a technology program. It is a strategic reinvention of the way insurance work gets done and, increasingly, it involves a choice about where to compete. Carriers able to spread AI infrastructure across a large premium base will compete on scale, while others will need to excel narrowly and partner for the capabilities they do not build.

The leaders we work with treat it as such and tackle it with the energy and discipline that a transformation demands. The early results are already visible. The window is open. The question for boards and CEOs is no longer whether to act, but how decisively.

[18] "How AI will reshape the economics of insurance: A CEO's guide to strategy," McKinsey, July 23, 2026.

Abhilash Sridharan is a partner in McKinsey's Mumbai office, where **Kamlesh Asrani** is an associate partner; **Alex Kimura** is a partner in the Singapore office, where **Azam Mohammad** is a senior partner; **Angat Sandhu** is a partner in the Sydney office; **Bernhard Kotanko** is a senior partner in the Hong Kong office; **Jan Henrich** is a senior partner in the Tokyo office; **Sid Kamath** is a partner in the Charlotte office; **Ritesh Agarwal** is an associate partner in the Kolkata office; and **Natasha Bhide** is a consultant in the Bengaluru office.

The authors wish to thank Abhishi Bhatia, Abhishek Shirali, Akash Lal, Alok Kshirsagar, Ben Ma, Brad Mendelson, Cezara Lisman, Charlene Wu, Diego Mattone, Fernando Serrano, John Qu, Kenji Nonaka, Koji Tazawa, Matthew Maloney, Michael Park, Mint Namasondhi, Peeyush Dalmia, Renny Thomas, Sebastian Kohls, Seungheon Song, Sudhanyu Veldurthy, Violet Chung, and William Bakara for their contributions to this report.

The authors also wish to thank colleagues in McKinsey's Insurance Practice, Operations Practice, Financial Services Practice, and QuantumBlack, AI by McKinsey.

McKinsey & Company
November 2026
Copyright © 2026 McKinsey & Company. All rights reserved.

McKinsey.com

@McKinsey
@McKinsey
@McKinsey

This publication is not intended to be used as the basis for trading in the shares of any company or for undertaking any other complex or significant financial transaction without consulting appropriate professional advisers.

No part of this publication may be copied or redistributed in any form without the prior written consent of McKinsey & Company.