Ultimate AI Strategy for Insurance

A White Paper by Artem Gonchakov

Simplifai

Executive Summary

The Core Argument

AI strategy cannot be separated from business strategy in insurance. When AI initiatives sit in IT or Innovation departments, they produce pilots. When AI integrates strategic choices about which customers to serve and how to compete, it produces operating leverage that moves combined ratios.

What Makes This Moment Different

Agentic AI systems can now plan and execute complete insurance workflows - not just assist humans but complete the work under governance. This architectural shift, combined with regulatory frameworks that exist in both the EU (AI Act, EIOPA) and the US (NAIC Model Bulletin), makes production deployment feasible for the first time. The 15% of carriers seeing material impact share one pattern: they wired AI directly into strategic workflows with governance built in from day one.

The Framework

This paper presents a systematic approach using Roger Martin's Strategy Choice Cascade as the foundation. Strategic assessment determines readiness across four dimensions (strategy, workflow, data, governance). Partner selection uses evidence-based evaluation, not vendor demos. Change leadership establishes clear metrics, approved risk boundaries, and delegated execution authority. First wins prove the model on bounded segments with production discipline. Value management ties outcomes to P&L with finance-grade attribution. Scaling follows concentric rings with expansion funding the next.

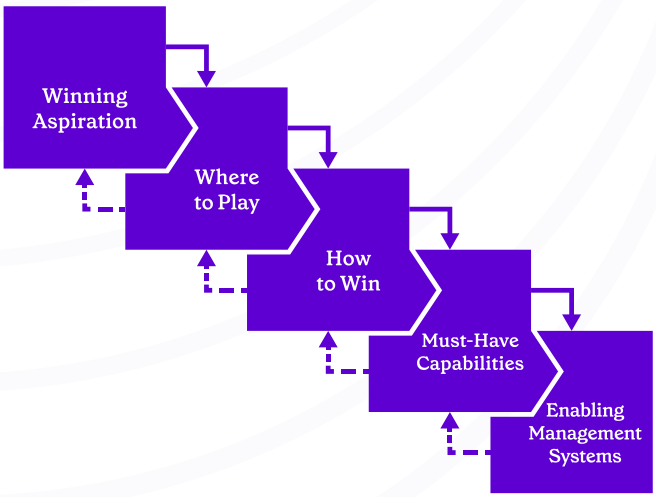

The Evidence

Carriers executing this approach report 30-40% productivity improvements, 25-35% cycle-time reductions, measurable loss-ratio gains, and audit-ready governance - all within 18-24 months from start to reaching 40-50% of eligible volume under automation. The gap between these leaders and carriers still running pilots is widening as premium growth slows and combined ratios approach 100%.

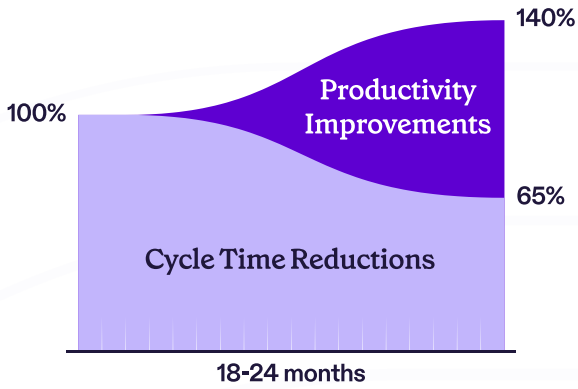

The Call to Action

Convene your executive team. Restate where you're playing and how you'll win. Identify 2-3 workflows where better decisions at scale create the most value. Baseline current performance. Set 90-day success criteria with measurable thresholds. Appoint a business owner. Start. The organizations that move now while the gap is bridgeable will pull away from those waiting for clarity that never arrives.

The organizations that move now while the gap is bridgeable will pull away from those waiting for clarity that never arrives.

Introduction

In 2025, the US and European insurance industry generated ~**$3.3 trillion** in premiums. Yet with combined ratios hovering at **99%**, the economics are stark: insurers pay out roughly **70 cents** of every premium dollar in claims and loss adjustment expenses, spend another **29 cents** on business operations, and keep just 1 cent as underwriting profit. The real profit? It comes from investing those premium dollars between collection and payout. This isn't the business model insurers advertise - it's the one they actually run.

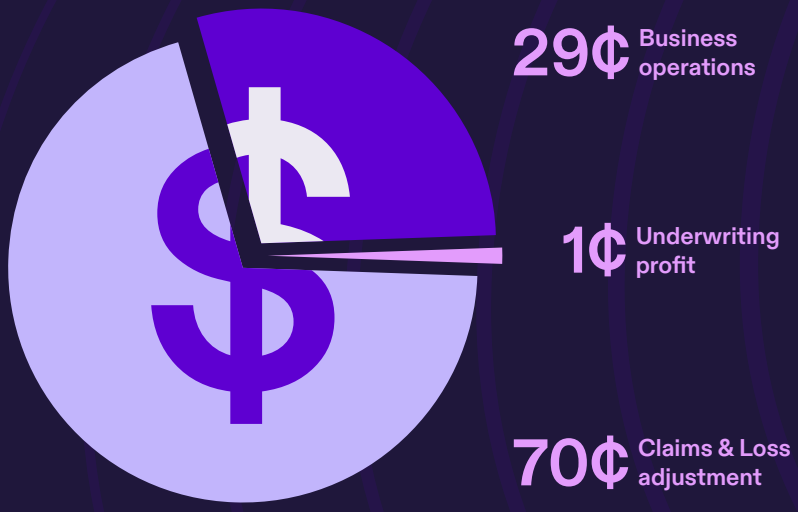

The pressure is mounting from every direction. Natural catastrophes are increasing in frequency and severity. Fraud schemes grow more sophisticated. Cyber risks multiply. Regulatory scrutiny intensifies. Customer expectations rise while loyalty falls. Meanwhile, traditional competitors and Insurtech startups alike are automating their way to lower expense ratios and faster service. The math is simple: carriers that can't reduce their cost structure while improving customer experience will lose both price and quality.

Insurance doesn't have an AI problem. It has a strategy execution problem. Technology works. The business model doesn't.

Most insurers respond with "digital transformation" - a sprawling initiative list where every department gets funding, no project gets killed, and 18 months later the metrics haven't moved. Marketing gets a chatbot. Claims get an image recognition tool. Underwriting gets a pricing model. IT gets a data lake. Nothing connects. The expense ratio climbs. Cycle times stay flat. A more focused competitor starts taking share.

AI has become the centerpiece of these digital transformation efforts, appearing in nearly every initiative. Carriers are testing machine learning for fraud detection, natural language processing for claims intake, computer vision for damage assessment, and generative AI for customer service. Technology works in controlled pilots. Investment is substantial - **83% of insurers** now allocate at least **$5 million** annually to AI, with **14%** spending over **$50 million**. Yet the pattern persists: promising demos, scattered production deployments, and minimal impact on the combined ratio.

The root problem isn't technology. It's treating AI as a collection of point solutions rather than a strategic capability. When AI sits fragmented across departments - each with its own vendor, its own data pipeline, its own evaluation criteria - it produces isolated improvements that don't compound. When AI integrates with business strategy - the fundamental choices about which customers to serve, which risks to accept, and how to compete - it produces operating leverage: underwriters make better decisions faster, claims handlers resolve cases with less friction, fraud analysts catch more with fewer false positives, and every process leaves an audit trail that satisfies regulators.

This paper shows how to make AI strategy part of your business strategy.

Specifically, it demonstrates how Agentic AI - systems that can plan, reason, and execute multi-step workflows under governance - addresses insurance's operational reality. Unlike earlier AI waves that require humans to interpret outputs, Agentic AI can orchestrate complete workflows from first notice of loss through payment, from submission through binding, from fraud alert through investigation. Technology has matured. The regulatory framework exists. The question now is execution.

What follows is a practical method for converting strategic choices into measurable AI capability. We start with what strategy actually means (and why most carriers confuse it with planning). Then we explain why Agentic AI specifically fits insurance operations. The core sections provide a phased approach: assessing readiness, selecting partners, leading change, delivering initial wins, managing value, and scaling safely. We address security, compliance, and ethics throughout - not as afterthoughts but as design requirements. We close with what it means to truly have an AI strategy versus simply having AI projects.

The window for competitive advantage is narrowing. Premium growth is slowing even as loss costs rise. The carriers that move decisively on production-grade, workflow-integrated AI will build compounding advantages in cycle time, loss ratio, and customer satisfaction that late movers cannot overcome. Some insurers are already there - processing claims in minutes instead of weeks, underwriting risks with data their competitors can't access, detecting fraud patterns that manual review would miss.

Let's start with the foundation - what business strategy actually is, and where AI fits within it.

What is an AI Strategy?

Most carriers have AI projects. Few have an AI strategy. The difference determines whether AI produces scattered productivity gains or compounding competitive advantages.

An AI strategy is the translation of business choices into a governed operating system that reliably produces faster, more consistent, explainable decisions at scale. It answers five questions that mirror strategic planning itself:

1. **What advantage are we pursuing?** Not "use AI" but "become category-best at X." For example: achieve the fastest claims settlement in personal property insurance or deliver the most accurate cyber risk pricing for SMEs. The advantage must be specific, measurable, and valuable to customers.
2. **Where will we compete for that advantage?** Which segments, products, channels, and geographies? An AI strategy that tries to improve everything improves nothing. Choose the ground where better decisions create the most value - perhaps broker-led SME cyber in Germany and Netherlands first, then Belgium next.
3. **How will AI enable us to win?** This is the technical architecture: Agentic AI systems that plan multi-step workflows, access the right data and tools, and operate under governance that regulators and boards can audit. For claims, this might mean FNOL intake that assembles files, checks coverage, proposes liability positions, and routes exceptions - all while citing sources and requiring human approval for material determinations.
4. **What capabilities must we build?** The data foundation: policy graphs, retrieval indexes over product rules, labeled evaluation sets for key workflows. The platform: orchestration, tool access, monitoring, and controls. The operating model: clear ownership, finance-grade measurement, and continuous improvement processes. Most importantly - the organizational muscle to wire AI into workflows rather than bolt it onto the side.
5. **How will we sustain this over time?** Management systems that fund based on measured outcomes, reallocate quarterly to what's working, and kill what isn't. Governance frameworks that satisfy EU AI Act, EIOPA principles, and NAIC guidance without slowing deployment. Evaluation harnesses that catch quality drift before it reaches customers. Unit economics that improve as volume scales.

The litmus test for boards and executive teams. Can your organization clearly state:

- Which business outcomes will AI improve, by how much, and by when?
- Which AI agents are in production, and what human gates do they pass through?
- What evidence shows that loss ratio, cycle time, or fraud leakage has improved versus baseline?
- How do you demonstrate data residency, explainability, and audit readiness to regulators?
- What happens if your primary AI vendor disappears tomorrow - do you own the critical assets?

Clear answers mean you have an AI strategy. Vague answers mean you have AI projects. The distinction matters because projects compete for budget while strategies shape the business. Projects produce local wins that don't compound. Strategy produces systematic advantages that competitors must match or lose market share.

The distinction matters because projects compete for budget while strategies shape the business. Projects produce local wins that don't compound. Strategy produces systematic advantages that competitors must match or lose market share.

Pilots are where strategy goes to die. They consume budgets, occupy talent, and produce PowerPoint decks. Strategy produces competitive advantages.

What makes this possible now - and urgent now - is the maturation of Agentic AI. Earlier AI generations (rules engines, machine learning, even generative AI) required humans to interpret outputs and take action. Agentic systems can do the work: assemble files, check rules, call APIs, propose decisions, route exceptions, and log everything for audit. This architectural shift is what finally makes "AI as business strategy" practical for regulated industries like insurance.

The rest of this paper shows exactly how to build and execute that strategy. But first, we need a common language for what business strategy itself means - because that's where AI strategy must begin.

Business Strategy

The Right Definition

In insurance, “strategy” gets confused with strategic planning more often than in most industries. Planning answers how, when, and by whom work gets done. Strategy is something harder: a falsifiable theory of competitive advantage over a multi-year horizon. It explains why this carrier will win, what trade-offs require, and how much uncertainty the organization can tolerate while testing the theory.

The distinction matters because the costs are real. Consider a composite carrier that spends nine months producing a “digital roadmap” with 120 initiatives - ensuring every function gets funding, and nothing gets stopped. 12 months later, the expense ratio is up, claims cycle time is flat, and a focused rival is taking market share. This pattern repeated across insurance for many years: fragmented, project-based transformation consistently underperforms coordinated; customer-back plays on both growth and efficiency.

Strategy is an integrated set of choices that must reinforce each other. Roger Martin’s Strategy Choice Cascade provides the clearest framework: a winning aspiration, where to play, how to win, the capabilities required, and the management systems that sustain them. Each choice constrains the next. Each involves real trade-offs.

The Five Strategic Choices

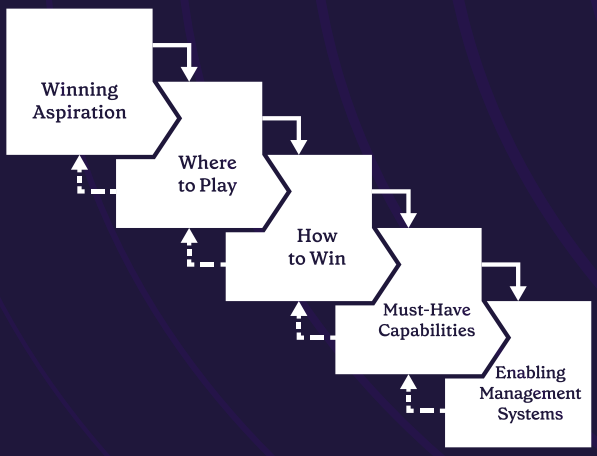

Winning Aspiration:

Articulate the ultimate goal in customer-centric terms, not just financial metrics. For an insurer, this might mean “become the go-to provider for climate-resilient property coverage in coastal markets” rather than “achieve 10% revenue growth.” The aspiration should be specific enough that you’d know whether you achieved it.

Where to Play:

Select specific markets, segments, channels, or geographies while making deliberate trade-offs. An insurer might choose small business cyber insurance in Germany and Netherlands, distributed through MGAs, while explicitly forgoing personal lines to avoid overextension. Every “yes” requires saying “no” elsewhere - otherwise it’s not a choice, it’s a wish list.

How to Win:

Define the unique value proposition and competitive edge. The answer must be specific and defensible. For cyber insurance, this might be “superior predictive underwriting using real-time threat intelligence from policyholders’ security tools, offering 20% better loss ratios than competitors using static questionnaires.” The test: why can’t competitors copy this? Perhaps they lack partnerships with security vendors, or the data science capability, or the willingness to invest ahead of proof.

Must-Have Capabilities:

Build distinctive skills and assets to enable the “how to win.” For the cyber example above, this requires advanced threat intelligence platforms, partnerships with security vendors, data science teams who understand both insurance and cybersecurity, and a culture that can balance innovation with regulatory compliance. This is where AI typically fits - as a core capability that amplifies competitive edges when aligned with strategy, not as a separate initiative pursued in isolation.

Management Systems:

Implement metrics, processes, and incentives to sustain execution. For the cyber play, this might involve dashboards tracking threat detection accuracy and claims frequency versus industry, with executive bonuses tied to these strategic outcomes rather than short-term sales volumes. If incentives pull one direction and strategy pulls another, incentives win every time.

Two Guardrails to Keep Strategy Honest

First.
Treat each box as a choice with a trade-off, not a wish list. If everything is a priority, nothing is a priority. The carrier above explicitly said "no" to personal lines. That freed capital, focus, and talent for the SME segments where they believed they could win.

Second.
Pressure-test the logic with "what would have to be true?" For each strategic choice, articulate the assumptions that must be held. For example: "German SMEs value speed over price" or "our AI can achieve 30% cycle time reduction with current data quality." These become testable hypotheses rather than hopeful statements.

**Strategy is expensive.
Everything is more expensive. Choose one.**

Where AI Fits

Once choices are explicit, AI's role becomes concrete: deliver decision advantage at the speed, scale, and accountability that "how to win" demands. The volumes of geospatial data, IoT feeds, and cyber telemetry now exceed human processing capacity. Customers and regulators expect decisions in hours, not weeks. And frameworks like the EU AI Act and NAIC Model Bulletin require documentation, explainability, human oversight, and audit trails that manual processes can't consistently provide.

AI enables the strategy in three specific ways:

1. **SPEED:** Agentic underwriting systems that pull external telemetry, check appetite rules, draft endorsements, and book risks - turning a 5-day process into 24 hours.
2. **ACCURACY:** Claims agents that ingest FNOL, verify coverage, assess damages using imagery and repair databases, detect fraud patterns, and route to the right adjuster - improving both cycle time and loss ratio.
3. **EVIDENCE:** Every decision documented with sources cited, rules applied, tools called, and approvals obtained - satisfying EIOPA's 2025 governance expectations without adding manual audit work.

Without the cascade, AI produces random automation. With the cascade, AI becomes the execution engine for strategic choices - turning "how to win" into measurable operating leverage.

Alignment is Where Strategies Live or Die

The cascade must flow through all organizational layers: corporate, business unit, functional, and operational. Large groups demonstrate how digital transformation sticks when these layers move together. When they don't - when underwriting modernizes while claims lags - the combined ratio doesn't budge, and the strategy fails in execution.

Regional or mono-line carriers can compress these layers, but the line of sight from aspiration to daily workflow still needs to be explicit. The underwriter evaluating a cyber risk should understand how their decisions connect to the "24-hour quote-to-bind" strategic commitment. The claims handler should see how their speed and accuracy affect the "sub-48-hour property settlement" promise. When that connection breaks, strategy becomes a document that sits in SharePoint while operations do what it's always done.

This is why AI strategy cannot sit in IT or Innovation. The cascade - from winning aspirations down to operational workflows - must be unified. AI is not a separate strategy. It's how you execute your business strategy at the speed and scale of modern insurance demands.

Now that we've established what strategy means and where AI fits within it, let's examine why Agentic AI specifically is the architecture insurance operations require.

AI is an evolving system

Roger Martin's Choice Cascade clarifies why and where to compete. AI clarifies how to execute at speed and scale that manual operations cannot match. But AI is not one technology - it's a stack of evolving techniques that only creates advantages when wired to specific choices about customers, products, channels, and controls.

In insurance, that wiring pays off in three places: sharper risk signals, faster and more consistent decisions, and evidence you can show to boards, auditors, and supervisors.

The Evolution: From Rules to Reasoning

This isn't a new ground for carriers. The industry has been automating decisions for decades, with each wave solving problems the previous generation couldn't handle:

- **Rule-based systems (1990s-2000s):**
  Automated eligibility checks, rating calculations, and simple underwriting decisions using "if-then" logic. Fast and consistent, but brittle - every exception required new code, and rules multiplied until systems became unmaintainable.
- **Classical machine learning (2000s-2010s):**
  Shifted from static rules to learned patterns in data. Pricing models found rating factors humans missed. Fraud detection caught patterns that rule sets couldn't be encoded. The limitation: required clean, structured data and couldn't handle text, images, or complex reasoning.
- **Deep learning (2010s-2020s):**
  Added the ability to learn from unstructured data - images for damage assessment, telematics for driver behavior, text for sentiment analysis. Transformed underwriting and claims by making sense of data types that previous systems ignored.
- **Generative AI (2022-2024):**
  Industrialized language work at enterprise scale - drafting policy summaries, generating claim notes, answering customer questions, and extracting information from documents. Independent analysts estimated GenAI's potential value creation in the trillions, with disproportionate impact where unstructured text dominates operations.

Each wave built on the last. Each solved a real problem. But each also had a fundamental limitation: they required humans to interpret outputs and take action. The model could score a risk, but an underwriter had to decide whether to quote. The system could flag a suspicious claim, but an adjuster had to investigate. The chatbot could draft a response, but someone had to review and send it.

Agentic AI: From Outputs to Outcomes

Agentic AI is the shift that converts productivity into operating leverage. Where GenAI produces content, agentic systems plan, reason, and act across multi-step workflows by calling tools, accessing data, and making decisions under explicit guardrails.

The architecture that works: Models plus tools plus governance. The model provides reasoning. Tools provide action - APIs to policy admin systems, payment platforms, data services, email, document generation. Governance provides control - approval requirements, human-in-the-loop gates, logging, and limits on what agents can and cannot do.

For insurance, this looks like:

- Claims intake that doesn't just extract information from FNOL - it assembles the complete file, calls policy admin to verify coverage, checks deductibles and limits, attaches relevant documents, assigns severity, proposes a handling path, and routes to the right team. All while citing sources and requiring supervisor approval for payments above threshold.

- Underwriting that doesn't just score a submission - it pulls external data (credit, loss history, property characteristics, weather exposure), checks appetite rules, identifies missing information and requests it, drafts endorsements for non-standard risks, calculates premium, and books the policy if it falls within delegated authority. Everything documented with the rules and data that drove each decision.

- Fraud investigation that doesn't just flag suspicious claims - it assembles evidence from multiple sources (claim history, social media, public records, vendor networks), identifies patterns, drafts investigation plans, schedules interviews, and produces investigation reports. All while maintaining a chain of custody for potential legal proceedings.

The difference is architectural. Previous AI generations made humans faster at their jobs. Agentic AI does the job with humans providing oversight, handling exceptions, and making judgments that require discretion or empathy.

Why This Matters Now

Three forces converge to make Agentic AI essential rather than optional:

VOLUME: The data inputs to insurance decisions now exceed human processing capacity. Underwriters evaluating property risks need to synthesize climate models, building characteristics, loss of history, neighborhood crime data, and local contractor availability. Claims handlers need to evaluate damage photos, repair estimates, coverage documents, and fraud indicators. Doing this manually for thousands of cases means either hiring proportionally (impossible at current margins) or making lower-quality decisions (unacceptable at current competition levels).

SPEED: Customer expectations and competitive dynamics have compressed decision windows from weeks to hours. Insurtech startups advertise instant quotes and same-day claims settlement. Traditional carriers responding with "we'll get back to you in 3-5 business days" lose the sale. But speed without accuracy destroys the loss ratio. Agentic systems can deliver both - processing in minutes while applying the same rigorous checks that took days manually.

GOVERNANCE: Regulators are raising the bar on explainability, oversight, and documentation. The EU AI Act requires documentation of data sources, testing results, human oversight mechanisms, and audit trails for high-risk AI applications. EIOPA's 2025 guidance expects insurers to explain how AI systems make decisions, demonstrate ongoing monitoring, and maintain records of model performance. The NAIC AI Model Bulletin emphasizes lifecycle controls and accountability. Agentic patterns that embed provenance, explainability, approval of workflows, and logging directly into the workflow are simpler to evidence supervisors than ad-hoc automations or purely manual processes.

The math is brutal: you can't hire fast enough to process the data. You can't slow down enough to stay competitive. Agentic AI isn't a choice anymore - it's the only way the arithmetic works.

Two Practical Tactics That Make This Work

Retrieval-augmented generation:

Ground AI outputs in your actual policy wordings, underwriting guidelines, claims procedures, and regulatory requirements. Instead of hoping the model "knows" your products, give it access to your authoritative documents and have it cited sources for every conclusion. This increases accuracy and creates the audit trail regulators expect.

Secure agent design patterns:

Scope what each agent can access and do. Require human approval for material decisions. Log every step with timestamps and attribution. Build evaluation harnesses that test against known-good outcomes. Create kill switches for immediate rollback if quality degrades. These aren't nice-to-have - they're the patterns that let agents run in production under the controls your organization already lives by.

The Bottom Line

GenAI handles language. Agentic AI handles work. When agents are wired into governed workflows, grounded in enterprise-specific knowledge, instrumented with the right tools, and bounded by approval requirements - you get decision latency measured in minutes instead of days, standardized execution that doesn't vary by who's working the case, and expertise that scales without linear headcount growth.

That's exactly what "how to win" demands in insurance. It's also why the next sections focus on what the data says, how to build this capability, and how to integrate it into operations under controls that boards and regulators trust.

What Data Says

Trends & Insights

The insurance industry's relationship with AI has reached an inflection point. The numbers tell two conflicting stories: massive investment meets modest impact. Understanding why reveals exactly where strategy separates winners from the pack.

The Investment Surge

AI spending in insurance has reached critical mass:

- **99%** of insurers have GenAI initiatives in scope or underway
- **83%** allocate at least **$5 million** annually to AI programs
- **14%** spend over **$50 million** per year
- Global AI in insurance market: **$10-11 billion** in 2025, growing at **30%+ annually**

The money is real. The commitment is there. So why isn't the combined ratio moving?

The Deployment Gap

Despite the investment, production deployment remains the exception:

- ~42% have deployed GenAI in at least one function
- But fewer than 15% report material impact on combined ratio, cycle time, or loss ratio
- Most common deployments: Customer service chatbots (limited claims impact) and document summarization (useful but not transformative)
- Least common: End-to-end workflow automation in underwriting or claims

The pattern is consistent across carriers: lots of pilots, limited production, minimal P&L impact. One European composite ran 47 AI experiments in 2024. Only three made it to production. None moved their expense ratio.

Where Leaders Are Pulling Ahead

A smaller group is seeing different results. Carriers that deploy Agentic AI into complete workflows - not just assistance tools - report:

- 30-40% productivity gains in claims and underwriting operations
- 25-35% cycle time reduction from FNOL to payment on automated-eligible claims
- 10–15-point improvement in straight-through processing rates
- Measurable retention uplift when cycle time and transparency improve

The difference isn't technological access. Every carrier can buy the same models and platforms. The difference is approach: workflow-first deployment with governance built in versus model-first pilots with integration as an afterthought.

The Three Barriers Holding Most Carriers Back

Industry surveys consistently identify the same obstacles:

1. **ROI Skepticism** (cited by 67% of carriers): Finance teams see pilot costs but struggle to measure benefits. “We saved 30 minutes per claim” doesn’t automatically translate to headcount reduction or margin improvement. Without attribution models that tie AI to P&L outcomes, pilots stay pilots.
2. **Governance Ambiguity** (cited by 58% of carriers): Risk and compliance teams face evolving regulations (EU AI Act, EIOPA guidance, NAIC bulletins) without clear playbooks. Rather than risk getting it wrong, many organizations slow or stop deployment. The irony: governed agentic patterns that embed controls actually make compliance easier, not harder.
3. **Integration Complexity** (cited by 54% of carriers): Policy admin systems from the 1990s. Claims platforms that don’t talk to billing. Data locked in departmental silos. AI that can’t access the right data at the right time can’t make good decisions. Carriers that solve integration first - building APIs, modernizing data access, establishing identity and consent frameworks - deploy faster and see better results.

The Macro Context Makes This Urgent

While carriers debate pilots, the business environment is tightening:

- **Premium growth slowing:** Global growth easing from ~5% in 2024 toward ~2% in 2025
- **Loss severity rising:** Natural catastrophe losses project to approach $145 billion in 2025
- **Combined ratios under pressure:** Industry average trending from 97% toward 99%+
- **Competition intensifying:** Both traditional competitors and insurtech startups automating operations

The window for experimentation is closing. Carriers that spent 2023-2024 on GenAI pilots now face a choice: commit to production-grade Agentic AI or watch competitors compound advantages in cycle time, loss ratio, and customer satisfaction.

What Separates Winners from Waiters

The data reveals a clear pattern. Carriers seeing material results share four characteristics:

1. They start workflow first: Not “where can we use AI?” but “which workflows create the most value if we can make them 10x faster and more consistent?” Then they architect AI to handle those workflows end-to-end.
2. They embed governance from day one: Provenance, explainability, human-in-the-loop gates, and audit trails aren’t added after pilots succeed. They’re requirements before the first line of code. This makes regulatory approval faster, not slower.
3. They measure operators: Finance owns the baselines. Every agent has a shadow P&L. Attribution uses holdout cohorts or difference-in-differences. Quarterly reviews reallocate capital to what’s working and kill what isn’t.
4. They package for reuse: The first claims workflow takes 90 days. The second takes 30 days. The fifth takes a week. Winners build component libraries - connectors, prompts, guardrails, policy graphs, evaluation sets - so each new deployment costs a fraction of the

The Bottom Line

The AI adoption gap is widening. A small group of carriers is deploying Agentic AI into production workflows, seeing 30-40% productivity gains and material improvement in strategic metrics. The majority remain stuck in pilot mode, accumulating costs without capturing value.

Technology isn't a bottleneck. Neither is regulatory clarity - frameworks exist in both the EU and the US. The bottleneck is approach: treating AI as a collection of point solutions instead of an integrated capability that executes business strategy.

The next sections show exactly how to cross that gap - from strategic assessment through partner selection, change management, first wins, value capture, and scaling under control. The playbook exists. The question is whether your organization will use it while competitive advantage is still available.

The carriers stuck in pilot mode aren't careful. They're slow. And in insurance, slowing has become expensive.

Building & Integrating AI Strategy

Converting AI ambition into operating leverage requires a sequenced approach that embeds governance from the start. Over the past 18 months, carriers that cleared the pilot-to-production gap followed a consistent pattern: they aligned to business outcomes first, designed controls before writing code, and treated the first deployment as miniature production rather than an experiment.

This chapter provides that playbook. It translates strategic choices into measurable AI capability through six phases: strategic assessment and readiness, partner selection, change leadership, delivering first wins, value management, and scaling safely. Each phase builds on the last. Each has clear entry and exit criteria. Together they turn scattered AI projects into a systematic competitive advantage.

Strategic Assessment and AI Readiness

Start with alignment, not technology.

Before evaluating technology or hiring data scientists, clarify the strategic anchor. Which business outcomes must AI improve, by how much, and by when? The carriers stuck in pilot mode are the ones that started with "let's try AI" instead of "let's move submission-to-bind time from 5 days to 24 hours for our target segment."

The foundation is alignment.

Restate your cascade choices - where you're playing and how you'll win - then identify the 3-5 workflows where better decisions at scale would create the most value. In practice, this means underwriting file assembly, FNOL triage, fraud screening, claims adjudication - the high-volume, multi-step processes where manual handling creates bottlenecks and inconsistency.

Assessment follows four dimensions

strategy clarity (cascade explicit, KPIs defined, trade-offs documented), workflow readiness (processes mapped, baselines measured, decision rights clear), data and technology foundation (systems accessible, documents digitized, logging in place), and organizational capability (business owners identified, risk embedded, evaluation methods defined).

The business case must speak the finance language.

Model cycle-time compression, labor redeployment, quality uplift, and experience improvement - each with dollar values tied to current baselines. Cost lines include platform fees, build effort, change management, and ongoing operations. Expect 3–5-year ROI windows for scaled programs, with faster payback where throughput is the primary lever. Carriers that articulate value this way clear investment committees. Carriers that tout "AI innovation" without P&L impact get budget cuts when priorities shift.

Governance design precedes code.

Classify target workflows under regulatory risk frameworks, then decide the guardrails: which decisions require human approval, what thresholds trigger escalation, how explainability will work, what gets logged, and for how long, where data must reside. Agents designed with these controls from day one reach production faster than agents where compliance is retrofitted after pilots succeed.

Choosing an AI Partner

The build-versus-buy decision is a trap. It's the wrong question, asked by organizations still thinking about AI as software rather than capability.

Here's the reframe. You're not buying an AI or building an AI. You're acquiring the organizational muscle to deploy, govern, and compound AI advantage faster than competitors. That muscle includes technology, yes - but more importantly, it includes pattern libraries, evaluation frameworks, and the institutional knowledge to ship workflow changes under governance in weeks instead of quarters.

Most carriers spend six months debating build versus buy, then spend eighteen months building or implementing, then discover they own neither the capability nor the speed they needed. The winning move is different: identify what you must own strategically (policy graphs, evaluation sets, orchestration patterns, customer data), then ruthlessly rent or partner for everything else.

The question isn't build or buy. It's: what do you own when the partnership ends? If the answer is 'not much,' you didn't buy a platform - you rented a dependency.

Requirements must map to regulated insurance operations.

The platform must orchestrate multi-step workflows with tool access and human-in-the-loop gates (not just generate text), provide provenance and explainability at every step (not just outputs without reasoning), operate under security and privacy controls that satisfy EU and US expectations (residency, isolation, audit trails), support governance frameworks aligned to regulatory requirements (EU AI Act, EIOPA, NAIC), and maintain open architecture (exportable artifacts, model neutrality, clear exit paths).

Selection should be evidence-driven, not demo-driven.

Establish a weighted scorecard: business outcomes (verified impact in comparable environments, finance-owned metrics, referenceable results), agentic capability (workflow orchestration, reliable tool-calling, graceful failure handling), governance and audit (provenance, evaluation harnesses, approval workflows), integration and data (API maturity, retrieval quality, identity handling), and security and compliance (certifications, residency, incident response). Then run a structured bake-off on your data with your evaluation criteria, not vendor demos on curated examples.

Commercial terms determine whether pilots can scale.

Structure pricing around workflow economics (per-case, per-task) is not opaque for token fees. Secure IP ownership for prompts, evaluation sets, and orchestration patterns. Prohibit vendor training on your data without consent. Require export rights for logs, indices, and artifacts. Set clear SLAs for uptime, latency, and incident response. Carriers that negotiate these terms up front avoid renegotiation of bottlenecks at scale. Carriers that defer commercial discussions until after technical validation often discover they've built strategic dependency into the architecture.

Partnership structure matters as much as technology selection.

Staff this as an operating capability, not a project: business owner, product owner, engineering lead, embedded risk partner on your side; solutions architect, governance specialist, dedicated support on the vendor side. Establish pattern libraries - reusable connectors, prompts, guardrails, evaluation sets - so each new agent costs 40-60% less than the previous one. Organizations that codify the reuse of compound velocity. Organizations that rebuild from scratch each time stay in linear scaling mode.

Leading the Change

Agentic AI reshapes work, decisions, and accountability. This transformation only sticks when leadership sets clear ambition, approves explicit boundaries, and delegates execution authority within those boundaries. Without top-down clarity, initiatives bounce between committees while competitors ship to production.

The CEO or business unit leader owns three commitments:

he metrics that must move (specific, measurable, time-bound), the investment envelope (budget, headcount, threshold for re-approval), and the governance boundary (what requires board approval versus operational delegation). These three commitments - outcomes, resources, authority - separate organizations that execute from organizations that pilot indefinitely.

Most CEOs ask: 'How do we govern AI?' The right question is: 'Who has authority to ship?' Governance without velocity is just an expensive hesitation.

Boards move faster when asked to approve risk envelopes, not feature lists.

Present the strategic linkage (which cascade choices this enables), target outcomes and ROI, control stack (decision thresholds, approval requirements, explainability methods), architecture and residency (one-page data flow), runbook and kill-switch (how we stop if quality degrades), and commercial terms (unit economics, exit rights). The board approves the risk class and delegation. The COO and CRO execute within approved bounds. This structure prevents boards from becoming bottlenecks while maintaining appropriate oversight.

Operating rhythm makes a change routine not special.

Weekly value stand-ups focus on coverage, quality, speed, safety, and blockers - outcomes and obstacles, no status of theater. Monthly governance reviews with audit and risk examine evaluation results, incidents, control effectiveness, and drift patterns - catching problems before they reach customers. Quarterly strategic reallocations fund what's working, fix or kill what isn't, and assess progress toward enterprise metrics. This cadence forces decisions and prevents indefinite drifts.

Adoption requires role-specific clarity, not generic training.

Claims adjusters and underwriters need to understand when to trust agent recommendations versus override, how to document their reasoning, and what happens to their workload. Supervisors need dashboards showing quality patterns and clear escalation paths. Auditors need to see complete decision trails. Executives need metrics to demonstrate business impact. Generic "AI awareness" sessions don't build trust. Specific answers to "what changes for me?" do.

In European markets, engage works of councils and data protection officers at design time, not review time.

Present AI as workflow transformation with redeployment plans, not headcount reduction. Complete DPIAs before pilots start. Demonstrate data minimization, consent mechanisms, and residency controls by design. Organizations that treat these as end-stage checkboxes face months of delays. Organizations that incorporate feedback early move faster and get better outcomes.

The pattern is consistent

leadership that sets clear direction and removes obstacles, governance that approves boundaries and delegates execution, operations that measure what matters and reallocate based on evidence. This organizational model - not the technology itself - determines whether AI produces isolated productivity gains or compounding competitive advantage.

Delivering the 1st Success

Pilots that look like miniature production succeed. Pilots that look like experiments stay experiments. The difference is treating the first deployment is proof of operating leverage, not proof of technology.

Claims is the natural starting point - text-heavy, tightly measured, and heavily governed. The journey from FNOL to payment offers clear metrics that boards and finance already track. More importantly, it demonstrates whether AI can operate under the controls of regulators' demand and customers expect.

The Strategic Approach

Start bounded, not broad. The first agent should handle one clean segment - perhaps property claims under a specific threshold with straightforward liability. Not because the technology can't handle more, but because proving the model on a controlled scope builds organizational confidence faster than attempting everything at once.

Design for audit from day one. Every decision must cite sources - which policy clauses, which data points, which rules applied. Every step must log - who approved what, when, and why. This isn't compliance theater. It's the architecture that satisfies EU AI Act transparency requirements and NAIC governance expectations while making the agent's reasoning inspectable by the humans who must trust it.

Set measurable gates before building. Define success with precision: 25-30% cycle-time reduction, 5-10 point accuracy improvement, under 2% unsupported recommendations, 100% audit completeness on material decisions. These thresholds separate meaningful progress from marginal tinkering.

The Pattern That Works

The agents that reach production share a common architecture: they're grounded in carrier-specific truth (policy wordings, underwriting rules, regulatory requirements), they orchestrate multi-step workflows with clear error handling, they call the right tools at the right time (policy admin, payments, third-party data), and they operate under explicit guardrails (approval thresholds, human-in-the-loop gates, kill-switch conditions).

This "models plus tools plus governance" pattern isn't theoretical. Carriers deploying this way report 30-40% productivity gains, 25-35% cycle-time reductions, and measurable improvements in straight-through processing - all while maintaining audit trails that satisfy regulators.

From Intake to Impact

The progression is deliberate. Start with intake - the agent ingests FNOL, verifies coverage, assesses severity, routes to the appropriate team. Once that operates reliably for 90+ days, add processing - document assembly, liability assessment, loss valuation, fraud triage. Then introduce payment under strict controls - approval matrices by amount, dual authorization above thresholds, complete audit packages. Finally close the loop with subrogation, reserve management, and continuous learning.

Each stage proves the model before expanding. Each stage packages reusable components - connectors, prompts, guardrails, evaluation sets - so the next deployment costs a fraction of the first.

What Measurement Reveals

The metrics that matter track speed (FNOL-to-decision time, straight-through rate), quality (accuracy versus human standard, override patterns, fraud yield), value (leakage avoided, labor redeployed), and governance (audit completeness, control effectiveness). Monthly reporting in one-page format keeps executives and boards focused on outcomes, not technology details.

Carriers that measure this way discover something important: the value isn't just operational efficiency. Faster cycle times improve customer satisfaction. Better accuracy improves loss ratios. Complete audit trails reduce regulatory friction. The compounding effects across these dimensions is what creates sustainable advantage.

The Discipline Required

Production agents demand production discipline. Evaluation harnesses with gold-labeled test sets. Automated regression testing on every change. Real-time monitoring with automatic rollback if quality degrades. Incident management with post-mortems and prevention measures. This operational rigor is what separates the 15% seeing material impact from the 85% stuck in pilots.

Scaling Systematically

Expansion follows a ring pattern: intake across all claims, then end-to-end for clean paths, then adjacent perils, then adjacent lines, then adjacent regions. Each ring opens only after the previous proves quality for 60+ days. This systematic approach - not opportunistic chasing of the next use case - is what allows carriers to reach 40-50% of eligible volume under governed automation within 18-24 months.

The first success does more than move metrics. It establishes the organizational muscle to wire AI into workflows, proves the governance model satisfies regulators, and demonstrates that strategy can execute at machine speed with human oversight. That capability - not the agent itself - is what compounds into competitive advantage.

What AI Agent can do across every claim

From First Notice of Loss to Final Closure - configured to your operations, running inside your ecosystem

STAGE 1

CLAIM INTAKE

From First Notice of Loss to a complete Claim File ready for investigation and with no missing information.

90%
of all admin work automated

75%
faster claim handling

- Capture the Loss (FNOL)
- Collect Evidence & Missing Information
- Verify Policy Coverage
- Categorize & Triage Claim
- Assign the Claim to Handler (or Team)
- Identify Large Losses
- Identify Customer Complaints
- Register new Claim
- Update existing Claim
- Create simple Claim Summary
- Policyholder Communications

STAGE 2

CLAIM PROCESSING

Investigate, assess, and decide on the claim with less leakage and fraud checks.

50%
reduced leakage

20%
less fraud

- Assess Liabilities
- Assess Damage and Repair Cost
- Calculate Reserve
- Identify & Pursue Recourse
- Generate Handler Action Items
- Run Leakage Checks
- Identify Fraud
- Make Decision or Recommendation
- Service Provider Communication
- Policyholder Communications

STAGE 3

CLAIM CLOSURE

Settle, recover, and close. Payment, quality, and compliance due diligence - every claim is compliant.

25%
lower cost per claim

100%
compliance ready

- Handle Payment
- Run Subrogation
- Run Recoveries
- Archive Case File
- Perform Quality Check
- Perform Compliance Check
- Close the Claim
- Policyholder Communications

Value Management

Expansion follows a ring pattern: intake across all claims, then end-to-end for clean paths, then adjacent perils, then adjacent lines, then adjacent regions. Each ring opens only after the previous proves quality for 60+ days. This systematic approach - not opportunistic chasing of the next use case - is what allows carriers to reach 40-50% of eligible volume under governed automation within 18-24 months.

The first success does more than move metrics. It establishes the organizational muscle to wire AI into workflows, proves the governance model satisfies regulators, and demonstrates that strategy can execute at machine speed with human oversight. That capability - not the agent itself - is what compounds into competitive advantage.

Value flows through five levers the P&L already recognizes:

Efficiency (fewer manual handoffs, reduced vendor dependencies), speed (shorter decision latency, compressed cycle times), accuracy (leakage avoided, fraud detected, subrogation recovered), scalability (more volume without proportional headcount growth), and transparency (decisions that regulators and customers can understand without extensive manual reconstruction). These levers are familiar. What's new is the magnitude and speed at which AI moves them.

Measurement must be defensible.

Treat each agent deployment like a discrete business unit with a shadow P&L. Revenue proxies include leakage avoided, fraud recovered, and subrogation collected. Operating costs include platform fees, infrastructure, support labor, and post-automation loss adjustment expenses. Finance owns the formulas. Operations owns the inputs. Monthly reconciliation prevents "productivity theater" - the appearance of progress without actual P&L impact.

Attribution separates signal from noise.

Use holdout cohorts or staggered rollouts to measure difference-in-differences against baseline performance adjusted for seasonality. Weight KPIs leaders already track and reconcile them monthly with finance. The carriers showing durable 30-40% productivity gains and measurable margin improvement share a common pattern: rigorous baselining and continuous evaluation, not one-time estimates projected forward.

Value leaks through predictable channels.

Adoption gaps occur when teams revert to manual processes. Process drift happens when workarounds bypass guardrails. Input quality degrades when data sources change without corresponding agent updates. Model decay emerges as product mix or customer behavior shifts. The controls that prevent leakage are straightforward: agent-first routing on eligible segments, change control on the golden path, input validation at ingestion, monthly evaluation against labeled test sets with automatic rollback if quality dips. These mechanisms align to governance expectations while keeping ROI credible over time.

Compounding value requires systematic reuse.

The first claims agent might deliver 30% productivity improvement and 18-month payback. The second agent in a different line should deploy in weeks, not months, by reusing 60-80% of components - connectors, retrieval patterns, approval workflows, evaluation frameworks. Organizations that package reusable assets see value compound: each new workflow becomes cheaper to automate, organizational muscle strengthens, and the gap versus competitors widens. Organizations that rebuild from scratch each time capture linear gains that don't accelerate.

The board and executive team should see value in their language:

which strategic metrics moved (submission-to-bind time, FNOL-to-payment time, loss ratio, expense ratio), by how much (with confidence intervals and attribution methodology), at what cost (all-in investment versus annualized benefit), with what trajectory (current run rate, projected path to target). One-page monthly reporting that finance has reviewed and signed keeps everyone focused on outcomes rather than technology enthusiasm.

Value management done well turns AI from a cost center into a competitive weapon - measurable advantage that accelerates rather than plateaus, financial returns that justify continued investment, and organizational capability that becomes harder for competitors to replicate. That capability is what scales.

Scale the Success

Once the first segment shows verified value and clean audit trails, the challenge shifts from “does it work?” to “can it work everywhere safely, repeatedly, and economically?” Scaling is less about adding features and more about systematizing how value is produced.

The scaling pattern is concentric rings, not scattered expansion.

Deploy first by claim amount (low to mid), then by peril (water to wind to theft), then by channel (direct to broker), then by region (home market to adjacent geographies). Each ring must pass quality and economic gates before the next opens: sustained volume at target metrics for 60+ days, zero P1 incidents, 100% audit completeness on material decisions, and finance-reconciled value that exceeds cost. This disciplined sequencing prevents the “deploy everywhere and hope” pattern that produces inconsistent results and erodes trust.

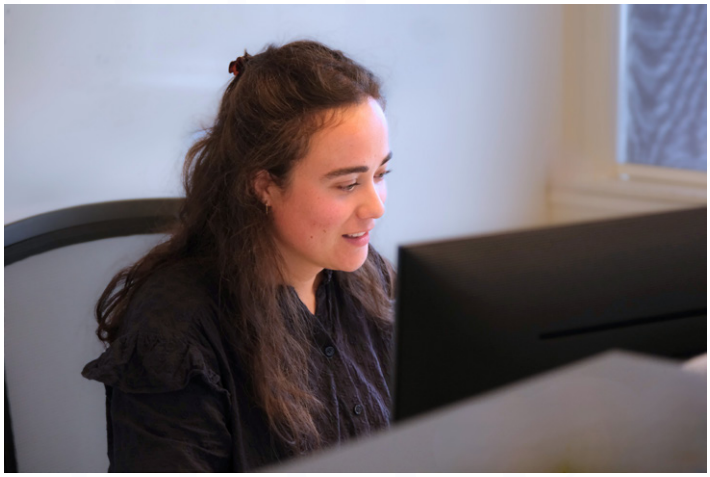

Reliability must scale alongside scope.

Establish service-level objectives that business operations recognize: decision quality versus expert standard, latency from input to decision, audit completeness, and tightly bounded unsupported-decision rates. Pair offline regression testing with online drift monitoring. Maintain one-click rollback capability. These practices mirror the engineering discipline required for any production system handling high-stakes decisions at volume.

Governance travels with every expansion.

The EU AI Act, EIOPA guidance, and NAIC Model Bulletin converge on the same expectations: documentation of data and decisions, testing before and after deployment, explainability that satisfies stakeholders, human oversight for material determinations, and records that survive audit. Treat these as design requirements, not compliance paperwork. Enforcement of least-privilege tool access, provenance at every step, dual authorization above thresholds, periodic red-team exercises, and appropriate data residency makes expansion read as “more of the same, safely” rather than new risk requiring extended review.

Unit economics must improve with scale, not degrade.

Track the real cost drivers: agent-task pricing, tool-call efficiency, cache and retrieval optimization, and loss adjustment expenses after vendor fees. Prefer pricing that matches workflows (per-case, per-decision) over opaque token-based models. Pool shared resources like embeddings and indices. Insist on export rights to preserve portability. The difference between programs that compound ROI and programs that stall at budget reviews often comes down to transparent unit economics and freedom to renegotiate or switch as leverage grows.

People determine whether scale sticks.

Run brief, role-specific enablement focused on when to trust, when to override, and how to document reasoning. Align incentives to the same outcomes leadership tracks - cycle time, accuracy, audit quality - so adoption is rewarded rather than optional. For parent company or board approvals, arrive with verified pilot results, controls mapped to regulations line-by-line, and phased rollout plans with clear kill-switches. This framing accelerates cross-organizational decisions and avoids prolonged committee reviews.

Set graduation criteria before opening each ring:

Sustained volume, targets met for consecutive weeks, zero critical incidents, 100% audit completeness, and finance-verified savings or recovery gains. Once criteria are met, open the next ring. Done systematically, scaling becomes a managed pipeline rather than a leap of faith. The same components that worked once now work across segments, lines, and regions - more work handled at lower cost, fewer errors, faster customer outcomes, and transparent records that satisfy oversight.

The carriers pulling away from competitors aren’t deploying fancier AI. They’re scaling systematically under control, compounding advantages with each expansion, and building organizational capability that becomes progressively harder to replicate.

Ensuring Security, Compliance, Regulations and Ethics

In insurance, AI must be explainable, controllable, and auditable. Regulatory expectations in the EU and US converge: the EU AI Act sets risk-based obligations, EIOPA emphasizes governance and transparency, the IAIS conducted thematic reviews on AI supervision, and the NAIC Model Bulletin stresses lifecycle controls and accountability. These aren't aspirational guidelines. They're the standards by which production deployments will be evaluated.

Control layers must travel with every agent and workflow.

Security is by design: zero-trust architecture, least-privilege access, tool allow-listing, encrypted data in transit and at rest, secret rotation, and network segmentation. Rate-limit outputs and red-team for agent-specific attacks like prompt injection, tool hijacking, and data leakage. Maintain kill-switch capability and one-click rollback to the last verified safe version.

Privacy and residency are non-negotiable in regulated markets.

Build for GDPR, UK GDPR, and US state privacy laws from the start: lawful bases for processing, consent where required, data minimization and pseudonymization, retention aligned to regulatory schedules, and EU policyholder data kept in-region with appropriate transfer mechanisms where needed. Identity verification and consent should be part of orchestration logic, not retrofitted afterward.

Model risk and explainability satisfy both internal and external stakeholders.

Each model and agent carries documentation: purpose, data sources, known limitations, approval status, and rollback path. Testing occurs on representative, labeled datasets - not vendor demos - with results segmented by relevant dimensions to surface fairness issues. Offline evaluation suites run continuously alongside online drift detection, feeding champion-challenger processes. Explanations must be human-readable: cite the policy clauses, data points, and reasoning steps behind each recommendation. This is what NAIC guidance and IAIS reviews expect.

Governance and audit make accountability explicit.

Define who owns business outcomes, who approves deployments, and who operates the platform. Require human-in-the-loop for decisions above materiality thresholds or when signals conflict. Log comprehensively: inputs, tool calls, intermediate reasoning, sources, decisions, approvers, and model versions. Maintain technical documentation, evaluation results, and post-deployment monitoring current for internal audit and supervisory review.

Third-party risk must flow down to vendors.

Contracts should limit data use (no training on customer data without consent), guarantee residency, require regular penetration testing, establish incident-response SLAs, secure export rights for logs and indices, and offer bring-your-own-key options. Favor architectures with model neutrality and clear data portability to preserve negotiating leverage and prevent lock-in.

Ethics and fairness go beyond technical accuracy.

Design for fair outcomes across customer segments. Run fairness assessments before and after deployment. Provide reason codes customers can understand. Ensure appeal mechanisms exist for adverse decisions, particularly in pricing, coverage, and claims. Align with consumer-duty principles in EU and UK markets and unfair-discrimination rules in US markets. These aren't compliance checkboxes - they're how trust is built and maintained with customers and regulators.

Handled this way, security, compliance, and ethics aren't obstacles to velocity. They're the architecture that lets AI run at scale under the controls boards trust, regulators expect, and customers deserve. Organizations that embed these from day one move faster than organizations that treat them as end-stage validation hurdles.

Conclusion

The thread running through this paper is straightforward: enterprise strategy sets the choices, AI executes those choices at speed and scale that manual operations cannot match. Treating AI as a side initiative produces pilots and press releases. Embedding it within the strategic cascade - the fundamental decisions about which customers to serve, which risks to accept, and how to compete - produces measurable shifts in the metrics that determine success or failure in insurance.

The insurance industry has always been good at managing risk. It's time to get good at taking it. The biggest risk isn't moving too fast on AI. It's moving too slow while calling it prudence.

The case for urgency is empirical, not speculative.

Premium growth is slowing from 5% toward 2%. Loss severity from natural catastrophes and cyber incidents continues rising. Combined ratios are trending from 97% toward 99% and beyond. Competition is intensifying from both traditional carriers automating operations and insurtech entrants built on modern architecture from day one. The carriers moving decisively on production-grade, workflow-integrated AI are building compounding advantages in cycle time, loss ratio, and customer satisfaction. The gap between leaders and laggards is widening, not narrowing.

The technology has matured past experimentation.

Agentic AI systems that plan multi-step workflows, access enterprise data and tools, and operate under governance frameworks are no longer theoretical. They're processing claims from FNOL to payment, underwriting submissions from intake to binding, and detecting fraud patterns that manual review would miss - all while maintaining audit trails that satisfy regulatory expectations. The bottleneck isn't capability. It's organizational approach: treating AI as scattered point solutions instead of integrated execution engines for business strategy.

The regulatory frameworks exist.

The EU AI Act provides risk-based requirements. EIOPA has issued governance principles. The NAIC has published model guidance. State insurance departments are actively supervising AI deployments.

Rather than creating barriers, these frameworks provide clarity: document your data and decisions, test before and after deployment, explain your reasoning, maintain human oversight for material determinations, and keep records that survive audit. Organizations that design to these requirements from day one reach production faster than organizations that treat compliance as validation theater at the end.

The path forward is systematic, not magical.

Assess readiness across strategy, workflow, data, and governance dimensions. Select partners based on proven outcomes in regulated environments, not demo quality. Lead change by setting clear metrics and delegating execution within approved risk boundaries. Deliver first wins on bounded segments with production discipline. Measure value in finance's language with attribution that survives scrutiny. Scale in concentric rings with each expansion paying for the next. Embed security, compliance, and ethics as design requirements, not afterthoughts.

The choice is binary.

Lead this transformation by wiring AI directly into strategic workflows under governance that boards and regulators trust - or watch competitors compound advantages while explaining to shareholders why your combined ratio trends in the wrong direction and your customer satisfaction lags the market. Some carriers are already there, processing in hours what competitors still measure in days, underwriting with data competitors can't access, detecting patterns competitors miss. The question isn't whether AI will reshape insurance operations. The question is whether your organization will lead that change or be disrupted by it.

What to do Monday morning:

Convene your executive team. Restate your cascade choices - where you're playing and how you'll win. Identify the 2-3 workflows where faster, more consistent decisions would create the most value. Baseline current performance. Define what success looks like in 90 days with measurable thresholds. Appoint a business owner, allocate resources, approve the risk envelope, and start. Not another study. Not another committee. Start.

The window for competitive advantage is measured in quarters, not years. The organizations that move now while the gap is bridgeable will pull away from those that wait for perfect clarity that never arrives. Strategy demands choices. AI demands execution. Together, they produce the operating leverage that separates market leaders from the pack.

The playbook exists. The question is whether you'll use it.

About Simplifai

Simplifai

Simplifai is an insurance-focused AI company founded in 2017, specializing in Agentic AI for claims operations. The company's platform enables carriers to design, deploy, operate, and govern enterprise-grade AI agents that handle complete workflows - from first notice of loss through payment and closure - under the governance controls regulated insurers require.

What differentiates Simplifai is the unity of platform and product.

AI agents are the product - systems that do real insurance work. The platform is how those agents are built, improved, and governed at scale. This integrated approach shortens time to production value while maintaining consistent controls across lines, regions, and regulatory jurisdictions. Carriers keep ownership of strategic assets - policy graphs, evaluation sets, orchestration patterns - that make each subsequent agent deployment faster and more economical than the last.

Proof points include 50+ AI agents live across multiple countries,

ncluding deployments with global cross-border claims operations and US personal and commercial carriers. The platform supports auditable execution through step-level provenance, replayable decision logs, and versioned policies - aligning with expectations from boards, internal audit, and supervisors in both EU and US markets.

The engagement model reflects the systematic approach described in this paper:

programs start with bounded journeys (often low-severity claims), operate as miniature production with governance from day one, and scale in rings across perils, segments, and regions. Value tracking is finance-grade - cycle time, straight-through rates, leakage and fraud impact, loss adjustment expenses - while risk and compliance co-own the control stack covering explainability, approval workflows, data residency, and record-keeping requirements.

Each successful deployment is packaged into reusable pattern libraries, allowing subsequent agents to reach production in weeks rather than months. This systematic reuse - not one-off custom builds - is what enables carriers to reach 40-50% of eligible volume under governed automation within 18-24 months while maintaining the audit readiness and control effectiveness that regulated operations demand.

Simplifai offers carriers an end-to-end capability for making Agentic AI operational in insurance - combining product and platform so strategic choices can execute at decision speed, operational consistency, and regulatory accountability that move combined ratios and create sustainable competitive advantage.

For organizations ready to move from AI projects to AI strategy, Simplifai provides the capability to start Monday morning.

Executive FAQ

Why now - not six months from now?

The gap between leaders and laggards is widening. Premium growth is slowing from 5% to 2% while loss severity rises and combined ratios approach 100%. Carriers deploying production-grade Agentic AI today are compounding advantages in cycle time and loss ratio that late movers cannot recover. Waiting for perfect clarity means explaining to your board why competitors are processing claims in hours while you're still measuring in days.

What's the difference between Agentic AI and GenAI?

GenAI writes. Agentic AI works. GenAI produces content - summaries, drafts, extractions. Agentic AI orchestrates complete workflows - planning multi-step processes, calling tools and data sources, making decisions under governance, and leaving auditable trails. The architecture is models plus tools plus governance. For insurance, this means FNOL-to-payment or submission-to-bind without manual handoffs, not just better document summaries.

How do we avoid pilot purgatory?

Anchor every initiative to one business metric with a pre-agreed threshold. "Reduce FNOL-to-payment from 14 days to 3 days for sub-€10K property claims by Q4 2026." Not "improve customer experience" or "leverage AI innovation." Run the pilot as miniature production with governance from day one. Set 90-day decision gates: scale, pivot, or kill based on evidence. Programs stall when they start with technology curiosity instead of business outcomes and when controls are added after success instead of designed in from the start.

What's the ROI model?

Value flows through five levers: efficiency (fewer manual touches), speed (compressed cycle times), accuracy (leakage avoided, fraud caught, subrogation found), scalability (volume growth without proportional headcount), and transparency (audit-ready decisions). Model each with finance-owned baselines and attribution using holdout cohorts. Typical targets: 3-5 year payback for scaled programs, faster where throughput is the primary lever. Carriers showing durable gains measure continuously and reconcile monthly with finance, not project forward from pilot estimates.

Build, buy, or hybrid?

Keep compounding assets in-house: policy graphs, retrieval indices, evaluation sets, orchestration patterns. These make the second agent cheaper than the first and preserve strategic independence. Use platforms for speed, governance maturity, and proven patterns in regulated environments. The decision hinges on time-to-impact versus control. Highly differentiated workflows with proprietary logic may justify builds. Most carriers see faster production deployment and better risk management with platforms that embed governance from day one.

What changes for risk, compliance, and audit?

Nothing gets lost. Requirements become design inputs, not end-stage validation. EU AI Act transparency obligations, EIOPA governance principles, and NAIC Model Bulletin expectations all converge: document data lineage, test before and after deployment, explain reasoning, maintain human oversight for material decisions, keep replayable logs, and demonstrate appropriate residency controls. Agents built this way clear regulatory review faster than manual processes retrofitted with documentation because explainability and oversight are architectural, not procedural.

What about people and jobs?

Agents remove repetitive, low-judgment work. People focus on complex cases, disputed liability, customer relationships, and situations requiring empathy or discretion. Most carriers redeploy rather than reduce - claims handlers move from routine intake to complex negotiations, underwriters move from data entry to risk assessment and relationship management. Adoption improves when role redesign is explicit, upskilling is resourced, and incentives align to strategic outcomes rather than activity volume.

How do we manage vendor risk and lock-in?

Own the strategic assets. Require export rights for prompts, workflows, evaluation sets, logs, and indices. Insist on model neutrality so you're not tied to one LLM provider. Favor architectures with clear data portability and documented exit procedures. Negotiate these terms before signing, not when you need leverage. Carriers that codify ownership and portability up front preserve strategic options. Carriers that defer these discussions discover they've built vendor dependency into the foundation.

What about data privacy and residency?

Run in-region where regulations require it. Enforce consent and purpose limitation at design time. Redact PII at ingestion points. Align documentation - DPIAs for EU, records of processing, consent logs - to supervisory expectations before deployment. Platforms built for regulated industries handle these requirements as configuration, not customization. Platforms built for general enterprise use often require extensive professional services to achieve the same compliance posture.

What's a realistic timeline to value?

Ninety days to prove safety, value, and operability on one bounded segment with production discipline and governance. Two quarters to scale the proven pattern across adjacent segments or perils. Four quarters to reach 40-50% of eligible volume under governed automation while tracking ROI against finance-owned baselines. Organizations that compress this timeline usually skip foundational work and end up in pilot loops. Organizations that extend it are treating AI as a project instead of strategic capability.

How do we know if we're ready?

Ask five questions: Can we name which business outcomes AI will improve, by how much, and by when? Can we describe which agents would run in production and what human gates they'd pass through? Can we show evidence that loss ratio, cycle time, or fraud leakage has improved in pilots versus baseline? Can we demonstrate to regulators how we ensure data residency, explainability, and audit readiness? If our primary vendor disappeared, do we own the critical assets - indices, evaluation sets, orchestration patterns? Clear answers mean readiness. Vague answers mean foundation work remains.

References

Strategy and Business Management

1. Martin, R. L., & Lafley, A. G. (2013). Playing to Win: How Strategy Really Works. Harvard Business Review Press.
2. Partaker, E. (2025). A Plan Is Not a Strategy. Harvard Business Review. https://hbr.org/2025/05/a-plan-is-not-a-strategy
3. Mintzberg, H. Quoted in A-Z Quotes. https://www.azquotes.com/author/10183-Henry_Mintzberg

Insurance Industry Analysis and Outlook

4. Deloitte. (2025). 2025 Global Insurance Outlook. https://www.deloitte.com/us/en/insights/industry/financial-services/financial-services-industry-outlooks/insurance-industry-outlook.html
5. EY. (2025). Global Insurance Outlook 2025. https://www.ey.com/content/dam/ey-unified-site/ey-com/en-gl/insights/insurance/documents/ey-gl-global-insurance-outlook-01-2025.pdf
6. Swiss Re Institute. (2024). Sigma 5/2024: Global Economic and Insurance Market Outlook 2025-26. https://www.swissre.com/institute/research/sigma-research/sigma-2024-05-global-economic-insurance-outlook-growth-geopolitics.html
7. Swiss Re Institute. (2025). Sigma 2/2025: World Insurance in 2025. https://www.swissre.com/institute/research/sigma-research/sigma-2025-02-world-insurance-riskier-fragmented-world.html
8. Swiss Re Institute. (2025). US Property/Casualty Outlook January 2025. https://www.swissre.com/institute/research/sigma-research/Insurance-Monitoring/us-property-casualty-outlook-january-2025.html
9. Insurance Information Institute (III). (2025). Facts + Statistics: Industry Overview. https://www.iii.org/fact-statistic/facts-statistics-industry-overview
10. Insurance Information Institute (III). (2025). A Firm Foundation: How Insurance Supports the Economy. https://www.iii.org/publications/a-firm-foundation-how-insurance-supports-the-economy/introduction/insurance-industry-at-a-glance

US Regulatory and Industry Bodies

11. NAIC. (2024). Property/Casualty and Title 2024 Mid-Year Industry Report. https://content.naic.org/sites/default/files/pc-and-title-2024mid-year-industry-report.pdf
12. NAIC. (2023). AI Model Bulletin. https://content.naic.org/sites/default/files/cmte-h-big-data-artificial-intelligence-wg-ai-model-bulletin.pdf.pdf
13. U.S. Department of the Treasury. (2024). Federal Insurance Office Annual Report. https://home.treasury.gov/system/files/311/2024-09-30%20Clean%20FIO%20AR%20508_2.pdf

EU Regulatory Framework

14. European Union. (2024). Regulation (EU) 2024/1689 on Artificial Intelligence (EU AI Act). https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ%3AL_202401689

15. EIOPA. (2021). AI Governance Principles. https://www.eiopa.europa.eu/system/files/2021-06/eiopa-ai-governance-principles-june-2021.pdf
16. EIOPA. (2025). Opinion on AI in Insurance. https://www.eiopa.europa.eu/document/download/8953a482-e587-429c-b416-1e24765ab250_en?filename=EIOPA-BoS-25-007-AI+Opinion.pdf
17. IAIS. (2023). Regulation and Supervision of AI/ML: A Thematic Review. https://www.iais.org/uploads/2023/12/Regulation-and-supervision-of-AI-ML-a-thematic-review.pdf

AI in Insurance: Industry Research

18. McKinsey & Company. (2025). The Future of AI in the Insurance Industry. https://www.mckinsey.com/industries/financial-services/our-insights/the-future-of-ai-in-the-insurance-industry
19. McKinsey & Company. (2025). Reimagining Insurance with a Comprehensive Approach to Gen AI. https://www.mckinsey.com/industries/financial-services/our-insights/reimagining-insurance-with-a-comprehensive-approach-to-gen-ai
20. McKinsey & Company. (2023). The Economic Potential of Generative AI: The Next Productivity Frontier. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier
21. McKinsey & Company. (2025). Seizing the Agentic AI Advantage. https://www.mckinsey.com/~/media/mckinsey/business%20functions/quantumblack/our%20insights/seizing%20the%20agentic%20ai%20advantage/seizing-the-agentic-ai-advantage.pdf
22. EY. (2024). Generative AI in Insurance: From Exploration to Implementation. https://www.ey.com/content/dam/ey-unified-site/ey-com/en-gl/insights/insurance/documents/ey-gl-generative-ai-in-insurance-05-2024.pdf
23. Conning. (2025). 2025 Survey on AI & Insurance Technology. https://www.conning.com/about-us/news/ir-pr---ai-survey-2025
24. ServiceNow. (2025). Enterprise AI Maturity Index 2025. https://www.servicenow.com/content/dam/servicenow-assets/public/en-us/doc-type/resource-center/white-paper/wp-enterprise-ai-maturity-index-2025.pdf

Technical Guidance on Agentic AI

25. OpenAI. (n.d.). A Practical Guide to Building Agents. https://cdn.openai.com/business-guides-and-resources/a-practical-guide-to-building-agents.pdf
26. OpenAI. (2025). New Tools for Building Agents. https://openai.com/index/new-tools-for-building-agents/
27. OpenAI. (n.d.). Agents Guide. https://platform.openai.com/docs/guides/agents
28. Thomson Reuters. (2025). Agentic AI 101: What Your Business Needs to Know. https://www.thomsonreuters.com/en/reports/agentic-ai-101
29. Google Research. (2025). Google's Approach for Secure AI Agents. https://storage.googleapis.com/gweb-research2023-media/pubtools/1018686.pdf
30. Palo Alto Networks. (2025). AI Agents Are Here. So Are the Threats. https://unit42.paloaltonetworks.com/agentic-ai-threats/

Catastrophe and Market Data

31. Aon. (2025). 2025 Climate and Catastrophe Insight. https://assets.aon.com/-/media/files/aon/reports/2025/2025-climate-catastrophe-insight.pdf
32. Aon. (2025). 1H 2025 Global Catastrophe Recap. https://www.aon.com/getmedia/01e165ae-1788-4997-a51b-9225bce850dd/1H-2025-Global-Catastrophe-Recap.pdf
33. Aon. (2025). Q1 2025 Global Catastrophe Recap. https://assets.aon.com/-/media/files/aon/reports/2025/q1-2025-global-catastrophe-recap.pdf

Industry News and Analysis

34. Risk & Insurance. (2025). Insurance Industry Accelerates AI Technology Adoption. https://riskandinsurance.com/insurance-industry-accelerates-ai-technology-adoption/
35. Wolters Kluwer. (2025). 2025 Insurance Tech Trends: AI, Big Data and Cautious Adoption. https://www.wolterskluwer.com/en/expert-insights/2025-insurance-tech-trends-ai-big-data-and-cautious-adoption

Industry Practitioner Insights

36. Allianz. (2025). From Generative AI to Agentic AI. https://tech.allianz.com/en/who-we-are/tech-blog/article-detail-012.html
37. McKinsey. (n.d.). Allianz Direct: Advancing as Europe's Leading Digital Insurer. https://www.mckinsey.com/capabilities/mckinsey-digital/how-we-help-clients/rewired-in-action/allianz-direct-advancing-as-europes-leading-digital-insurer

Additional Resources

38. Statista. (2025). AI Adoption in Insurance Worldwide 2024-2025. https://www.statista.com/statistics/1610910/ai-adoption-insurance-worldwide/
39. Ataccama. (2025). Insurance AI Use Cases Whitepaper. https://www.ataccama.com/whitepaper/insurance-ai-use-cases
40. Vonage. (2025). AI in Insurance 2025: How Insurers Can Harness the Power of AI. https://www.vonage.com/resources/articles/ai-in-insurance/