

A WHITEPAPER SERIES FOR
INSURANCE CIOs AND CTOs

Simplifying Agentic AI for CIOs and CTOs

Part 1 of 3

sapiens 

Every Chief Information Officer (CIO) and Chief Technology Officer (CTO) struggle with the same fundamental question: where exactly does AI fit in your existing architecture?

This whitepaper series, **Bridging Worlds: Agentic AI and the Traditional Insurance Technology Stack**, cuts straight through noise to give you practical guidance on how modern insurance platforms can successfully integrate agentic AI through strategic partnerships with leading AI providers.

Executive Summary

This whitepaper outlines why AI implementation feels so complicated. This includes figuring out where agentic AI will fit in with the traditional insurance three-tier architecture. “Add-on-AI” won’t work, so we explain the new layer that facilitates a true agentic AI experience and what this means beyond the buzzwords.

You’ll also receive decisions you can act on the **three components that make AI agentic**, plus **two techniques that make AI truly intelligent**.

The Challenge Every Insurance CIO/CTO Now Faces

Walk into any insurance company today and you'll find a familiar three-tier architecture that's been the backbone of operations for decades:

- 1 Your **user interaction layer** handles everything, from agent portals to customer mobile apps.
- 2 The **business logic layer** contains all the insurance rules and workflows that took years to perfect.
- 3 And the **data layer** stores everything from policy details, to claims history.

Traditional Insurance Three-Tier Architecture




Figure 1: Insurers' traditional three-tier architecture – the established foundation that agentic AI enhances, rather than replaces.

This stack works. It's reliable. And it's auditable. So, when AI vendors start speaking about "transformation," most CIOs and CTOs I know have the same reaction: where exactly do these new capabilities fit and what happens to everything we've already built?

The statistics explain the confusion. While 84% of enterprise CIOs believe AI will be as significant as the internet, only 11% have implemented it fully. The gap exists due to a fundamental misunderstanding of how AI will integrate with existing technology and systems.

What's happening on the ground? 76% of insurance executives say their organizations have implemented generative AI in at least one function, but most of these are point solutions that don't fundamentally change operations. They're AI "Band-Aids" slapped on traditional processes, not integrated, architectural symbiosis.

Why “Add-On” AI Misses the Point

The confusion makes sense when you look at how most AI implementations are positioned. Vendors either suggest replacing your core systems entirely, which CIOs and CTOs are hesitant to do for obvious reasons, or to treat AI as just another application that plugs in somewhere. Both approaches miss what agentic AI enables: a new orchestration layer that works with your existing systems, rather than replacing them.

Four-Layer Business Architecture Model

How AI integrates with traditional insurance operations




Figure 2: The new Agentic AI orchestration layer.

Understanding Agentic AI: Beyond the Buzzwords

“Agentic AI” has become another buzzword that vendors throw around without explaining its actual meaning for insurance operations. Let’s get practical...

Traditional automation in insurance is like a sophisticated vending machine. Press the right buttons in the right sequence and you get predictable results. If a claim amount exceeds the designated

high-value threshold, route it to the senior adjuster. If a policy lapses, send the renewal notice. Clear rules, clear outcomes.

Agentic AI is fundamentally different.

It’s more like having an intelligent assistant who understands your goals and can figure out how to achieve them, even when the situation doesn’t match

any predetermined rule.

This is a critical difference in insurance, because so much of what we do involves judgment calls, incomplete information, and coordination across multiple systems. Traditional automation does an excellent job with straightforward cases, but struggles with everything else.

The Three Components That Make AI “Agentic”

Every agentic AI system has three core components working together. Understanding these helps clarify where the technology fits in your architecture:

1

The Model

It acts as the reasoning engine. In insurance terms, it’s what reads a broker’s email, understands what’s being requested, and decides how to respond appropriately. It doesn’t replace your business logic – it enhances your ability to process unstructured information that traditional systems can’t handle.

2

Instructions and Memory

This component provides the context and guidance the AI needs to behave appropriately. This includes your underwriting guidelines, regulatory requirements, customer service standards, and memory of previous interactions. Your institutional knowledge becomes the AI’s operating framework.


3

Tools and Data Access

They enable the AI to interact with your existing systems. These components might pull policy information from your core platform, retrieve credit scores from external databases, or access pricing engines. They connect agentic AI to your technology stack without replacing it.

The Three Building Blocks of Agentic AI

Core Components That Work Together



Key Insight: These three components work together continuously, with information flowing bi-directionally to create intelligent, context-aware responses and actions.

Figure 3: The three building blocks of Agentic AI – model, instructions/memory, and tools working together to create intelligent orchestration.



Here's a simple example of how they work together: a broker emails a submission for a Cambridge, Massachusetts (MA) restaurant. The model understands the submission details. Instructions guide the AI to check restaurant-specific risk factors and remember this broker's track record. Tools automatically pull property records, local crime data, fire department proximity, and similar business claims. The system compiles a preliminary assessment with recommendations. The difference from traditional workflows: nobody programmed exactly how to handle "restaurant submission in Cambridge, MA." The AI understood the goal and figured out the necessary steps.

Two Techniques That Make AI Truly Intelligent

What makes agentic AI particularly valuable for insurance are two capabilities that go well beyond simple automation.

Chain-of-thought Processing

Adaptive Problem Solving

Chain-of-thought Processing

Chain-of-thought processing lets the AI break down complex problems into logical steps, like how an experienced underwriter approaches a complex risk assessment. Instead of just giving an answer, the AI shows its reasoning:

Step 1

Identify the business type and key risk factors

Step 2

Check for any unusual circumstances or red flags

Step 3

Compare to similar risks in our portfolio

Step 4

Apply company underwriting guidelines

Step 5

Provide recommendation with reasoning

Scenario: a broker submits an application for a \$2.6 million commercial property (restaurant) in city center Cambridge, MA.

Traditional system response: "Recommend standard terms at \$5,800 annual premium."

Agentic AI response: an agentic AI system transforms traditional underwriting by autonomously conducting multi-layered risk analysis with transparent, auditable reasoning. Rather than simply applying predetermined rules, the AI agent independently evaluates each submission through structured decisionmaking that mirrors expert underwriter thinking.

In the image below, the AI agent systematically progresses through property analysis, identifying key risk indicators like the commercial kitchen's elevated fire exposure, while recognizing

mitigating factors, such as superior emergency response access in the city center location.

The agent then dynamically applies underwriting guidelines, calculating multiplicative risk adjustments, and determining that the \$2.6M property value triggers mandatory senior review protocols. Throughout this process, the AI provides complete reasoning transparency, showing how it weighted construction materials against occupancy risks, balanced location disadvantages with condition improvements, and arrived at the final premium calculation.

This agentic approach enables consistent, explainable decisions, while maintaining the nuanced judgment that complex commercial risks demand, ultimately delivering faster turnaround times – without sacrificing underwriting quality or regulatory compliance.





Step 1

Property Analysis

- Commercial restaurant, **\$2.6M** value, central Cambridge, MA
- High-foot-traffic area with good emergency services access (fire station nearby, strong municipal response)



Step 2

Risk Factor Assessment

- **Fire risk:** Commercial kitchen = elevated risk
- **Location risk:** Central Cambridge = moderate property crime, but strong fire/police protection
- **Construction:** Brick/steel = solid fire resistance
- **Occupancy:** Restaurant = higher liability exposure (slips/falls, food safety, etc.)



Step 3

Underwriting Guidelines Application

- Property value **\$2.6M** = requires senior underwriter review per guidelines
- Restaurant class = **apply 1.25x base rate multiplier**
- Central Cambridge location = apply **0.9x location credit**
- Recent renovations = apply **0.95x condition credit**



Step 4

Premium Calculation

- Base rate: **\$0.50 per \$100 of value** ($\approx$ \$2.76 per \$1,000)
Applied multipliers: $1.25 \times 0.9 \times 0.95 = 1.07$
- Final rate: **\$2.38 per \$1,000 of value**
- **Annual premium = \$5.660**

Figure 4: The agentic approach.

This matters enormously for insurance operations. Every decision can be traced and explained for audits. Junior staff can see expert reasoning patterns for training. Regulators can understand the decisionmaking process. Supervisors can review reasoning quality for oversight.

Adaptive Problem-Solving

Adaptive problem-solving means the AI can change its approach based on what it discovers, rather than following a fixed workflow.

If initial data suggests higher risk, it might automatically request additional information, or route to a senior underwriter.

Think about a beneficiary submitting a life insurance death claim where the medical records show some inconsistencies with the reported cause of death. Traditional systems apply a fixed claims investigation checklist.

Adaptive AI might start with standard verification, but then investigate

further based on specific medical inconsistencies it identifies, automatically requesting additional medical records, or routing to a specialist medical examiner. This is similar to how an experienced life claims investigator would approach the case.

Scenario: a beneficiary submits a death claim, but medical records show inconsistencies with the reported cause of death.

Traditional system response: Apply fixed fraud checklist

Agentic AI response: in life and annuities claims processing, agentic AI enhances

fraud detection and claims validation through autonomous investigative reasoning that goes beyond traditional rule-based systems, see **figure 5: Further investigation.**

The AI agent independently conducts claim analysis, cross-referencing multiple data sources to identify inconsistencies that human adjusters might miss. In the example below, the AI agent

demonstrates structured investigative thinking by first identifying medical discrepancies between death certificates and recent health records, then proactively seeking additional evidence through multiple API calls to build a complete picture.

Thought: This death claim has medical inconsistencies I need to investigate.

The cause of death doesn't align with recent medical records.

Action: MedicalRecordsAPI [policy_id: "LIFE-2024-8901"]

Observation: Recent cardiac test showed normal results, but death certificate lists "heart failure"

Thought: Medical records contradict cause of death. Let me verify the death certificate details.

Action: DeathCertificateAPI [certificate_id: "DC-2024-3456"]

Observation: Physician who signed certificate is not on hospital staff. Timing inconsistencies in records.

Thought: Invalid physician and timing issues suggest fraud. Check for recent policy changes.

Action: PolicyHistoryAPI [policy_id: "LIFE-2024-8901"]

Observation: Beneficiary changed 3 months ago, large policy loan taken recently

Thought: Multiple red flags indicate potential fraud. This needs SIU investigation.

Action: FlagForReview [claim_id: "LIFE-CLM-2024-9876", reason: "potential_fraud"]

Observation: Claim flagged for medical fraud investigation

Final Assessment: Death claim flagged for SIU review due to medical inconsistencies, invalid death certificate, and suspicious recent policy changes.

Figure 5: Further investigation

Rather than simply flagging predetermined risk indicators, the agent reasons through each discovery, recognizing that a non-hospital physician signing a death certificate combined with timing inconsistencies suggests potential document fraud. The AI then expands its investigation scope, uncovering

suspicious policy modifications and financial transactions that create a comprehensive fraud pattern.

This autonomous reasoning capability enables the agent to make practical judgments about when to escalate claims for special investigation unit review, providing detailed rationale for

each decision. The result is faster, more accurate fraud detection with complete audit trails, allowing legitimate claims to process quickly, while ensuring suspicious cases receive appropriate human oversight and investigation.

In summary, chain-of-thought processing and adaptive problem-solving make agentic AI systematic (like your best procedures) and flexible (like your best people).

Memory and Learning in Practice

Agentic AI systems maintain context in ways that traditional systems don't. Short-term memory keeps track of everything discussed in a single interaction – if a customer calls about a claim, the AI remembers the entire conversation. Long-term memory stores important patterns and preferences for future use.

This isn't just about storing data points. AI memory maintains context and continuity across interactions, which improves both customer experience and operational efficiency. The system learns what works and what doesn't, getting better over time, without requiring manual reprogramming.

Connecting to Your Data: Making AI Smarter

One of the biggest advantages of agentic AI is its ability to work with information that wasn't available when the system was originally built. This happens through Retrieval-Augmented Generation (RAG), which sounds complicated, but works simply.

The AI stores information from your documents, policies, and procedures in a searchable format. When a question arises, it automatically finds relevant information and combines it with the current context to provide accurate, up-to-date responses.

For example, a customer service agent asks about coverage for a specific type of damage. The AI instantly retrieves relevant policy language, recent bulletins, and similar claims examples to provide a comprehensive answer.

Training AI for Your Specific Needs

You can enhance AI performance for your insurance operations in three main ways:

1

Learning from Examples

This means showing the AI how to handle situations correctly. Demonstrate proper claims categorization and the AI learns your classification standards.

2

Access to Current Information

This occurs through RAG, which ensures responses are accurate and up to date with current policies and procedures.

3

Specialized Training

It's for highly specific tasks that can further train the AI on your unique data to internalize your business practices and decisionmaking patterns.

The practical approach: start with examples and external data access, then move to specialized training only if needed. Begin simply and add complexity based on results.

The Bottom Line: Agentic AI Makes Sense

Agentic AI delivers the most value when you need to handle complex decisions requiring judgment and context, maintain difficult-to-program rules that change frequently, process unstructured data like emails and documents, or coordinate processes spanning multiple systems.

For routine, rule-based tasks with clear procedures, traditional automation may still be more appropriate and easier to maintain.

Agentic AI isn't about replacing your existing systems. It means adding an intelligent orchestration layer that can understand and process unstructured information, coordinate activities across multiple systems, adapt to new situations without reprogramming, maintain context across complex workflows, and learn from examples to improve over time.

This understanding provides the foundation for seeing how agentic AI integrates with traditional insurance architecture to create more flexible, intelligent operations, while preserving the reliability and compliance controls that insurers need.



Next in This Series

Next, we'll get into the architectural specifics: how the evolution from three-tier to four-layer architecture actually works, integration patterns that connect AI with existing systems, cloud versus on-premise deployment considerations, and governance frameworks that maintain compliance, while enabling innovation.

Stay tuned for the rest of this white paper series:

- [Part 2: The architectural reality – how to evolve from three-tier to four-layer thinking](#)
- [Part 3: Real business scenarios and your practical next steps](#)

Contact Us

About Sapiens

Sapiens International Corporation N.V. is a global leader of AI-centric, SaaS-based insurance software, delivering hyper-relevant experiences that are efficient, compliant, and innovative. With agile intelligence, Sapiens' solutions turn real-time data and human insight into precise action at every moment, across every risk. The Sapiens platform includes agentic workflows accelerating every capability across policy, underwriting, claims, reinsurance, decisioning, and finance and compliance. With more than 600 insurers in over 30 countries running on Sapiens, our deep industry expertise is the foundation of our long-term relationships, from initial implementation through to modernization and market transformation. Sapiens is headquartered in London, serving customers in property and casualty, life, reinsurance, specialty, and workers' compensation from offices across North America, Europe, the Middle East, and Asia Pacific.

For more information visit sapiens.com or follow us on [LinkedIn](#).

© Copyright Sapiens International. All Rights Reserved. June 2026