GUIDE

Advance the Possible: How to Build an Autonomous Insurance Platform

Glen Ceniza, Chief Technology Officer, Sapiens

sapiens

Executive Summary

Global insurers spend an estimated 70% of their IT budgets maintaining legacy systems built for batch processing, manual workflows, and siloed data, while McKinsey estimates AI could unlock more than $1.1 trillion in annual value across the industry. The Sapiens Autonomous Insurance Platform (SapiensAIP) closes that gap with a single, reusable AI layer that sits on top of your existing core systems and connects underwriting, policy, billing, claims, and customer engagement into one environment where AI agents not only recommend tasks but actually perform them.

AIP is built on three layers: a persona-based experience layer for the people who use it, an intelligence layer where agentic flows do the work, and a foundation layer, including our data moat or “ontology” based on 40+ years of industry experience. It’s the “Brain” that gives every AI agent the context to reason like an insurer, not a generic model. It deploys independently of which core system you run and is designed to be adopted one use case at a time rather than all at once.

This playbook is written for the people asking the hard questions about SapiensAIP: what problem it solves, how the architecture works, why Sapiens is positioned to deliver it, and what your organisation should do to prepare.

CHAPTER 1

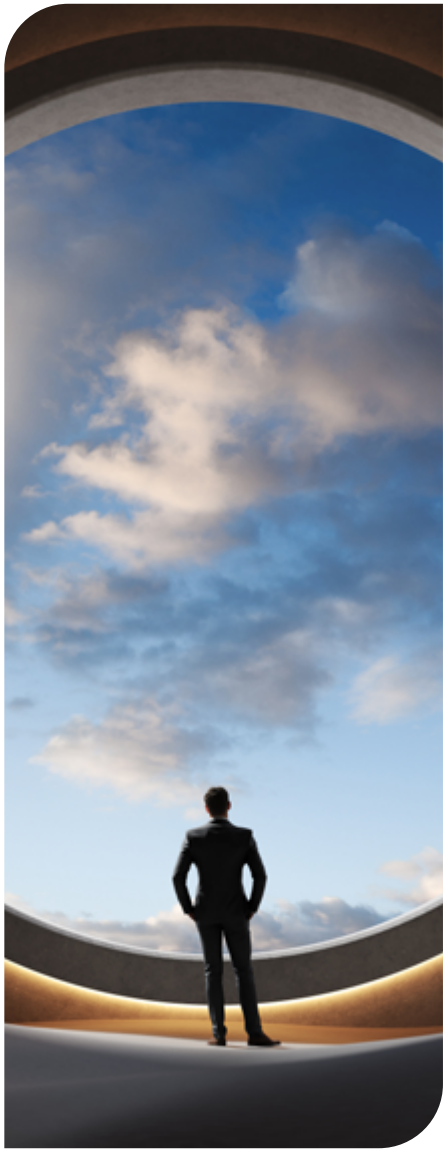
Technologists Are Planning for a Future That's Already Here

Most insurers operate on architecture built over decades, coordinating separate systems for policy, claims, and billing through point-to-point integration. That fragmentation is expensive: a survey of 350 senior insurance leaders found 35% are making business decisions based on siloed systems, and Accenture estimates poor claims experiences alone cost the industry upwards of $170 billion. For a technology leader, the real cost isn't the systems themselves; it's the data they trap, which is exactly what AI needs to work well.

Technical readiness varies enormously between insurers, and often between departments within the same organisation. Some are cloud-native and running production AI use cases; others remain on premise, particularly when policyholder Personally Identifiable Information (PII) and financial data make regulators and risk teams cautious. These issues persist even as the business need for a unified AI platform has already been agreed upon by key stakeholders.

It's hard to overlook the level of bluster in a market that's still inventing itself. "Agentic" and "autonomous" are used interchangeably, and the number of agents a vendor claims to run is frequently treated as a maturity signal. For us, an agent that converts a name from "William" to "Bill" is not the same as one that reads an entire claims file and produces a settlement estimate. Technical evaluators need a way to compare platforms on outcomes, not headlines.

As consumption models for agentic AI are also evolving, insurers are right to ask how a platform's pricing interacts with AI licenses for Claude, Copilot, and other frontier platforms they may already hold. Large insurers are also reasonably wary of concentrating their entire AI strategy in one vendor. Any serious evaluation needs a clear answer on both: how usage and cost are measured, and how the platform coexists with the rest of the technology you already have.



CHAPTER 2

Why Only an AI Platform Strategy Will Save You

Most insurers didn't set out to build AI confusion; it just happened one point solution at a time. A chatbot here, a fraud model there, an underwriting assistant bolted onto a legacy core. Each tool solves a narrow problem, but none of them talk to each other, and none of them see the whole customer. That's the real barrier to AI's promise in insurance: not a lack of models, but a lack of connection.

A platform strategy is the solution that's right under your nose. When agentic AI runs on one system that already understands policies, claims, underwriting, and compliance, insight travels instead of getting stuck. A claims signal informs pricing. An underwriting decision informs service. Data doesn't need to be re-explained to every new tool. If it is on the same platform, it's already in context, already governed, already trusted.

Point solutions can't do that work. They can automate a task. They can't reason across a book of business, adapt to a shifting regulation, or act with the judgment a real underwriter brings. Only a platform gives AI the full picture, and the guardrails, to act with confidence, not just speed.

We aren't trying to add more automation. A platform strategy gives insurers one foundation where agentic AI can learn continuously, act consistently, and stay compliant as rules and risks change. That's what makes AI hyper-relevant instead of merely impressive: it knows the customer, the context, and the moment, because it's built on the same platform that runs the business.

Insurers piecing together point solutions will keep hitting the same wall: smart tools, disconnected outcomes. Insurers who build on a single platform will get AI that compounds, gets sharper with every interaction, and moves the business forward together.

> The promise of AI in insurance was never about isolated tools. It's about a unified, autonomous platform, built for the next 30 years of this industry.

sapiens

CHAPTER 3

What an Insurance AI Platform Looks Like and Acts Like

SapiensAIP is built on a three-layer technology stack, deployed as a standalone, reusable AI layer on top of your existing core system. And, as we have been emphasizing, there needs to be a core system in place. Whether your policy administration runs on a Sapiens core or another vendor system, it doesn't matter, your platform plugs in without requiring you to re-platform first.

At the top is the experience layer: persona-based interfaces for customers, agents, brokers, and configurators. This is the only layer most users ever see.

Beneath it is the intelligence layer, where the agents live. Rather than a single agent replicating an entire job, a platform like Sapiens allows you to build agentic flows. These purpose-built agents are chained together to complete a function such as an underwriter's or claims handler's day-to-day work. Organisations can adopt a full flow if they are comfortable with a high degree of automation or adopt individual agents within it and keep more human quality gates in place.

The foundation layer is where the platform's real value sits: your agent studio where your agents are customised, your compliance and audit logging, and your "ontology" or data moat that gives every agent the reasoning and domain context it needs. At Sapiens, we call this layer "The Brain." The Brain was originally built as an internal engine to power Sapiens' own agentic flows; only through early customer engagements did it become clear how valuable it could be as a standalone capability. Sapiens is now working on how to expose more of it directly. Customers never interact with this layer directly, but it is what separates an agent built on a general-purpose model from one that understands underwriting, claims, and regulatory context out-of-the-box.

Because your platform is designed to be reusable across geographies and lines of business without depending on a specific core system, most of the data an out-of-the-box use case needs already lives in your existing core. For this reason, initial data preparation is typically light. Use cases that go beyond the standard library, such as a custom document-verification workflow or an integration with an internal ticketing system, need the relevant source data mapped and shared up front, but that is scoped per use case rather than required as a wholesale migration.

sapiens

CHAPTER 4

Why SapiensAIP Should Be Your Choice

SapiensAIP is built on more than 40 years of insurance domain logic, with more than 2,000 modelled workflows spanning property & casualty, life & annuities, reinsurance, and workers' compensation. Domain intelligence at this depth cannot be added after the fact. It is integrated into the platform itself, connected through interfaces such as MCP (Model Context Protocol) connectors, and trained on the operational reality of each insurance vertical, rather than on data that simply mentions insurance in passing. That distinction shapes everything about how SapiensAIP is built and brought to market as the most definitive answer to the challenges insurers face when evaluating their AI strategy.

Our Unique Value Proposition

1 Domain intelligence, not a generic model

Four decades of insurance-specific workflows and an ontology purpose-built for underwriting, claims, and policy administration mean AI agents start from insurance reasoning, not from a general-purpose model that has merely been pointed at insurance data. That difference shows up immediately in how these agents interpret a claim, price a risk, or flag an exception. This is the Brain that reasons the way an experienced underwriter or claims handler does, because that operational logic is what the platform is built on.

2 A platform, not a stack of pilots

Our AIP connects the full lifecycle, from underwriting, policy, billing, claims, and customer engagement through one reusable AI layer, rather than a set of siloed automations bolted onto separate systems. Siloed pilots tend to plateau at incremental efficiency gains because each one solves a narrow problem in isolation, with its own data model, its own integration effort, and its own maintenance burden. A single connected layer means the intelligence built for one workflow compounds into the next, instead of resetting at every departmental boundary. It also means the second and third use cases you deploy get cheaper and faster to stand up.

3 Adopt at your own pace

Agentic flows in SapiensAIP are built from individual agents, which means adoption isn't an all-or-nothing platform switch. You can take on a full flow end-to-end, or select specific agents and quality gates that match a narrower use case. This framework lets you calibrate automation to your organisation's comfort level and risk appetite, expanding scope only once you've proven value at a smaller scale.

4 Built-in observability

A dedicated cloud and observability layer tracks query volume and consumption from day one, giving technical buyers a real foundation for governance and cost forecasting rather than having to build that visibility themselves after the fact. This capability is currently being calibrated through Sapiens' beta programme, so expect it to mature further as the platform moves toward general availability.

5 A path for the AI you already run

Most insurers already have investments in AI licenses, and Sapiens is actively defining how SapiensAIP complements those tools rather than requiring you to abandon them. Work is also underway on how the platform fits alongside other vendors already in your tech stack, so adopting AIP doesn't mean unwinding decisions you've already made.

6 Design-partner validation, not a lab demo

SapiensAIP is being shaped by a live beta programme with customers spanning property & casualty and life, across North America, South Africa, the UK, and the Nordics. That geographic and line-of-business spread matters: an insurer in South Africa and an insurer in the Nordics operate under different regulatory regimes, different claims volumes, and different customer expectations, so a platform validated across all four regions is being pressure-tested against real variation, not a single friendly environment. What ships reflects direct feedback from insurers running real workflows in real regulatory environments, not a roadmap built in isolation from customer input.

7 A defined path for off and on-premise

For insurers keeping policyholder data on-premise for regulatory or risk reasons, moving to a SaaS-native AI platform can feel like a leap. Sapiens' SaaS upgrade pathway is built specifically to make that migration a deliberate, staged upgrade rather than a leap of faith, so the decision to modernise your infrastructure and the decision to adopt AIP don't have to happen in the same breath.

Taken together, these points describe a platform built the way insurance actually works: layered on decades of domain-specific logic, proven with live customers before general release, and designed to meet insurers wherever they are in their cloud journey rather than forcing a single path. That's a different proposition from a generic AI model retrofitted with insurance vocabulary, and it's the reason the groundwork described in getting your organisation ready is worth doing well before SapiensAIP reaches general availability.

How to Get Your Organization Ready

Technical readiness for SapiensAIP is less about a large upfront project and more about laying six pieces of groundwork before you begin. None of them require a transformation programme. Together, though, they determine how quickly you move from pilot to production, and how much value you see once you get there.



1 Map your core and cloud posture

Start by knowing, precisely, where you stand today: on-premises, hybrid, or SaaS-native. This isn't a prerequisite for adoption: an AIP deploys the same way regardless of which core system you run, but your cloud posture shapes the sequencing and timeline of your rollout. An organisation already SaaS-native will move through onboarding faster than one still managing on-premise infrastructure. Also, knowing which stage you're in lets you set realistic expectations from day one rather than discovering constraints mid-project.

2 Choose a first use case, not the whole platform

Resist the instinct to treat adoption as all-or-nothing. Identify one or two priority workflows (e.g., a claims document check, a specific step in underwriting) and prove value there first. A narrow, well-chosen starting point gives your teams a concrete, low-risk way to build confidence in the technology, generates early wins you can point to internally, and creates a template you can extend to the next workflow rather than a big-bang rollout with no natural checkpoint.

3 Stage the data behind that use case

The data burden depends entirely on what you're building. Standard use cases need little upfront data work, since an AIP reuses data already sitting in your core system. A bespoke use case is different: it needs its specific source data identified, cleaned where necessary, and made available in advance. Knowing which category your chosen use case falls into prevents data readiness from becoming a mid-project surprise.

4 Agree internally on what "autonomous" means to you

Before you sit down with us, discuss the best use of AI agents, agentic flow, and autonomous systems within your own team. Different platforms offer different levels of autonomy, and without shared internal goals, evaluation conversations drift toward counting features and agents rather than assessing outcomes. Getting alignment first lets you compare vendors on what their agents actually accomplish, not on how many they claim to have.

5 Redefine how you measure success

Feature counts are a poor proxy for value. Shift your evaluation criteria toward achieved business outcomes: settlement speed, loss ratio, risk-selection accuracy, retention. Just as importantly, get your business and technology stakeholders aligned on these metrics before the project starts.

sapiens

6 Bring your underwriting, claims, and operations leaders in early

These are the people whose judgement becomes the training signal for your agents. Involving them only at the review stage after a use case has been defined and built wastes their most valuable contribution. Bring them in at the definition stage instead, when their expertise can shape what the agent should do, not just whether it did it well once it's finished.

None of this groundwork is heavy lifting on its own. But skipping it tends to surface the same problems later, at a point where they're more expensive to fix. For example, you might encounter a use case that's too broad to prove out, data gaps discovered mid-build, or a vendor evaluation that never quite gets to the metrics that matter. Doing the groundwork first is what makes AIP adoption a fast, low-risk step rather than a leap.



Checklist for Getting Started

Readiness and platform fit are only useful if they lead somewhere. Once you've mapped your core and cloud posture, picked a first use case, and aligned your stakeholders on what success looks like, the next question is simply how to move from evaluation to a live pilot without losing the discipline you've just built. A formal evaluation of SapiensAIP typically follows these steps:

- **Start a scoping conversation**
  Map your core systems, priority use cases, and cloud posture against the AIP architecture with your Sapiens contact.
- **Ask about a design-partner engagement**
  If direct input into how your priority use case is built matters to you, ask whether an arrangement that is like the current beta cohort is available for your evaluation window.
- **Agree upon pilot scope and success metrics up front**
  Define the business outcomes you will measure before implementation work begins, not after.
- **Plan a phased rollout**
  Validate the foundation layer against your core and data, prove out one agentic flow end-to-end, then expand into further use cases and lines of business at your own pace.
- **Set governance early**
  Agree upon usage guardrails, spend thresholds, and observability requirements with Sapiens as part of the commercial conversation, not as an afterthought.
- **Clarify how AIP fits your existing AI investments**
  Make sure the evaluation accounts for your current AI licenses and vendors, not just what AIP replaces.
- **Schedule a technical architecture review**
  Once your first use case is defined, confirm data readiness and integration requirements before implementation begins.

Key takeaway

The insurers already running this way are not piloting AI; they are running their business on it. SapiensAIP is built to help you advance the possible and make that move on your own timeline, starting with the use case that solves a nagging business problem.

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, 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 modernisation 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, Asia Pacific and the Middle East including Israel where the company was founded.

For more information visit sapiens.com or follow us on LinkedIn.

© Copyright Sapiens International. All Rights Reserved. October 2026

sapiens