

What carriers should be building in 2027

2026 was the year the agent stopped being interesting. Two questions decide the next one: what you get to keep, and what you can prove.

WHITEPAPER

Author
Michael Robinson

ust.com



■ U
S T

Table of contents

• The bill nobody itemises	1
• Stop thinking about "a model"	2
• Two questions decide what goes in the weights	3
• The part almost nobody is selling	4
• What this does not mean	6
• Five things worth doing in the next six months	6
• The shape of the next few years	7

The bill nobody itemises

Walk any insurance technology conference this year, and you will hear the same sentence from every stand. We are deploying agents. It is true, it is expensive, and it has stopped meaning anything. **Agentic deployment has become an entry requirement, not a differentiator.** By 2028, the two things that will separate carriers are ownership and evidence.

When every carrier in a segment can do a thing, that thing is a cost of entry, not a strategy. The interesting question is no longer whether an insurer can put an agent into a claims process. It is what the insurer has at the end of three years of doing it, why any individual decision came out the way it did, and whether it can evidence both to a regulator, an auditor, an ombudsman, or a capacity provider.

Every carrier understands the first cost of enterprise AI. There is a second cost, and it never reaches the invoice. Every correction your adjusters make. Every prompt your underwriting team refines over eleven iterations until it finally handles the edge case properly. Every evaluation you run to satisfy your own risk function that an output is safe. All of that work improves this quarter. Reputable providers commit contractually not to train on your data, and that commitment matters. But it does not change the arithmetic. Nothing accumulates on your side.

Year three looks like year one, minus the spend.




Figure 1 — The second cost of renting. Illustrative shape, not measured data.

And because every carrier in your segment is renting from the same handful of providers, whatever capability you gain, the carrier across the street gains on identical terms, in the same quarter, for the same money. Differentiation that comes out of a shared model is not differentiation. It is a shared upgrade that everyone paid for separately.

Meanwhile, what actually separates you sits outside the model entirely. Your products. Your wordings. Your appetite. How your underwriters read a submission that is technically in appetite but feels wrong. How your adjusters weigh a file. What your organisation has learned, over decades, to treat as normal. A general-purpose model knows insurance. It does not know your insurance, and no amount of prompting teaches it permanently, because nothing you teach it survives the context window.

The uncomfortable framing

Most insurers do not own their policy administration, claims or billing platforms. They licence them, and a decade of cloud migration moved them further from ownership rather than closer. That was a reasonable trade: a system executes a process, and processes are broadly common across the industry. There was never much advantage in owning the code that issues a policy. A model is a different kind of thing. It does not execute your process. It reproduces your judgement, and your judgement is not available from any vendor, because no vendor has it.

Stop thinking about "a model"

The framing that unlocks this is to stop asking which model to buy and start asking which layers to own. There are four, and they behave completely differently.




Figure 2 — Four layers, and how differently they behave. Ownership is a question you answer per layer, not once.

The foundation layer is a choice, not a build; this is an open-weight model you can deploy and adapt. Worth noting early rather than late: you own the weights you train, but you inherit the licence of the base you started from, and those licences are not uniform. Some are permissive; others carry attribution, field-of-use, or scale conditions that flow through to everything derived from them. Choosing a base model is a legal decision at least as much as a technical one, and it belongs at the start of a project.

The industry layer teaches the language of insurance. It offers no competitive advantage, which is why it should be shared infrastructure rather than something each carrier funds alone. The company layer is where advantage lives. Your forms, your terminology, your underwriting rules, your claims philosophy, your operating standards.

And the agent layer is where today's facts arrive at runtime. An agent running on general-purpose intelligence starts every task as a stranger to your business and spends its context being briefed. An agent standing on the layers below already knows how your organisation reasons, and spends its context on the work.

Two questions decide what goes in the weights

This is where most sovereignty conversations go wrong. The instinct is to train everything you own. That produces something expensive and confidently out of date. They are the two questions from the cover, in working clothes: what you keep is what you train; what you must prove is what you wire.

How fast does it change? Your terminology, your house voice, your claims philosophy, and the way your organisation reasons about risk move on a scale of years. Train those. Your endorsement library, your rate tables, your form versions and your desk procedures move on a scale of weeks. Train those, and you have built something that will be confidently wrong by next quarter.

Does it need to be provable? Some knowledge stays outside the weights even when it is perfectly stable, because the answer has to cite a source. When an agent tells an insured what their policy covers, the regulator wants the form referenced. Retrieval is not only about freshness. It is about producing an audit trail.

Train the judgment. Wire in the facts. A model estate is not one large artefact; it is a stable trained core surrounded by machinery that keeps it current and accountable.




Figure 3 — The two-question sort. One quadrant earns the weights; everything else is wired in and cited.

The part almost nobody is selling

Here is the consequence that is missing from every sovereign-AI pitch.

Ownership moves the liability.

Under the EU framework, a firm that uses a vendor's AI system carries a relatively light duty set. A firm that adapts a model and puts the result into service under its own name looks a great deal more like a provider, and the provider duty set for a high-risk system is an order of magnitude heavier: risk management system, data governance covering your training data, technical documentation, logging, accuracy and robustness, a quality management system, conformity assessment, registration, and post-market monitoring for as long as the thing is in service.




Figure 4 — Indicative mapping of eu ai act high-risk obligations. The duties apply from 2 december 2027 — which makes 2027 the build window.

The specific classification of any given system is a legal question that turns on facts, and the high-risk obligations themselves do not bite until December 2027 following this summer's amendments. Take advice on your own position. But the commercial shape of the argument is not in doubt, and it cuts directly against the way sovereignty is currently being marketed: the more of your intelligence you own, the more of the regulatory apparatus you have to run yourself.

That is not an argument against ownership. It is an argument against ownership without machinery. A carrier that builds a company model and does not simultaneously build the control plane has bought an asset and an unfunded obligation at the same time.

The good news, if you are in insurance, is that you already know how to do this. You have run delegated authority for a century. An agent is a delegated authority. Its limits are a binder schedule. Its output is a bordereau. Its oversight is a coverholder audit. Its failure is a breach of authority. The control vocabulary does not need to be invented for AI. It needs to be extended to it. Extended, it takes five artefacts: an authority schedule saying what each agent may decide; a decision ledger recording what it did and why; provenance on every fact it relied on; a drift report against the authority granted; and an attestation pack your auditor, board and capacity provider can consume. All but one are descendants of documents your delegated authority team already runs.

An agent is a delegated authority

The control vocabulary does not need inventing for AI – it needs extending to it.

A century of delegated authority		Extended to the agent	
01	The coverholder	→	The agent
02	The binder schedule	→	Authority schedule Scope . limits . a named owner
03	The bordereau	→	Decision ledger Every decision, with reasons
04	The coverholder audit	→	Drift report Behaviour vs granted authority
05	Breach of authority	→	Revocation Suspend . unbind . replay

Figure 5 — The control vocabulary, extended. Only the provenance record has no paper ancestor.

What this does not mean

It does not mean own everything. A model you own sits at a fixed capability level while frontier models improve past it. You take on the evaluation burden, the retraining cadence, the drift monitoring, and the safety work that a frontier vendor would otherwise absorb on your behalf. For open-ended reasoning at the edge of what is possible, renting is the right answer and probably always will be. General intelligence is becoming steadily more accessible and steadily cheaper. Your institutional knowledge never will.

Ownership earns its keep on a specific profile of workload: the ones that run constantly, touch your most sensitive data, depend on vocabulary only you use, and separate you from the carrier across the street. That is exactly where renting compounds against you. Everywhere else, rent, and be glad of it.

Five things worth doing in the next six months

- 1  **Inventory what you are teaching and losing.** Count the corrections, the refined prompts, the evaluation sets your teams have built this year. That volume is the asset you are currently discarding, and quantifying it is the fastest way to make this conversation real to a CFO.
- 2  **Run the two-question sort on one process.** Take claims triage or submission clearance. Split every piece of knowledge it uses into "changes slowly and shapes judgement" and "changes fast or must cite a source".
- 3  **Test the performance claim before you believe it.** One constantly running task, one adapted model, one frontier baseline, one shared evaluation set, cost-adjusted.
- 4  **Get the classification question answered by counsel.** Would owning this make you a provider? What would that oblige you to run?
- 5  **Build the evidence machinery on the first agent, not the fortieth.** Authority schedules, decision ledgers, provenance records and drift reports are cheap to design in and painful to retrofit. The carriers whose agent programmes stalled this year did not have a model problem. They had an audit committee that could not get comfortable, which is the same problem arriving later and more expensively.

The shape of the next few years

The last decade of insurance technology was about owning less: licence the platform, move to cloud, consume the service. That was correct for systems. It breaks for models, because a model does not execute your process, it reproduces your judgement.

The carriers who come out of this well will not be the ones who deployed the most agents in 2026. They will be the ones who, by 2028, own the layer that reflects how they actually work, rent everything above and below it without embarrassment, and can produce on demand the evidence that every decision those layers made was authorised, bounded, monitored and reversible. The high-risk duties arrive in December 2027, and the machinery takes about a year to build well. That is what 2027 is for.

Ready to move beyond AI pilots?

Discover how leading insurers are operationalizing AI with stronger governance, auditable decision-making, and workflow integration.

Connect with a UST expert →



Together, we build for boundless impact

Since 1999, UST has worked side by side with the world's best companies to make a powerful impact through transformation. Powered by technology, inspired by people, and led by our purpose, we partner with our clients from design to operation. Our digital solutions, proprietary platforms, engineering, R&D, products, and innovation ecosystem turn core challenges into impactful, disruptive solutions. With deep industry knowledge and a future-ready mindset, we infuse expertise, innovation, and agility into our clients' organizations—delivering measurable value and positive lasting change for them, their customers, and communities around the world. Together, with 30,000+ employees in 35+ countries, we build for boundless impact—touching billions of lives in the process.

ust.com

© 2026 UST Global Inc.

Version 0102-20261016

**U ■
S T**