Rocket® software

The Insurance Industry's Next Competitive Advantage

Governed agentic AI at scale

Contents

Executive summary

Insurers are moving quickly on AI. Nearly every carrier has a pilot underway, whether in claims triage, underwriting support, customer service, or document processing. The early results look promising, and the pressure to move faster keeps building.

But speed is about to become table stakes. When most insurers can stand up an AI pilot relatively quickly, being fast no longer sets you apart. The harder and more valuable capability is scaling AI into core systems without creating regulatory, compliance, or operational risk. That's where the real separation will happen.

This paper makes a straightforward argument. In regulated insurance environments, governance won't slow you down. It's the thing that lets you move at all. The carriers who win the next phase won't be the ones who deployed AI first. They'll be the ones who can prove that every automated action stayed within policy, remained traceable, and respected existing controls, even as they scaled from one use case to dozens.

Rocket® EVA™, Rocket Software's governed agentic AI platform, is an example of a working example of what governed agentic AI looks like when it's built for core insurance operations rather than bolted on afterward.

The competitive pressure insurers face

The insurance industry is under real pressure to modernize, and AI has become the focal point. Carriers are watching operational costs climb, experienced staff retire, and customer expectations rise. AI promises to help on all three fronts, and boards are asking pointed questions about where their organization stands.

That pressure has produced a wave of experimentation. Teams are testing generative models to summarize claim notes, draft correspondence, answer coverage questions, and help underwriters process submissions faster. Much of this works. It saves time, impresses stakeholders, and builds appetite for more ambitious deployments.

Here's the problem with treating speed as the goal. When the barrier to entry is low, everyone clears it. An AI pilot that summarizes documents or drafts a letter is genuinely useful, but it's also something a competitor can replicate quickly. The advantage it creates is thin and temporary.

The durable advantage lies somewhere harder to reach. It's in the systems of record: the claims platforms, policy administration systems, and batch cycles that actually run the business. Moving AI into those environments, where an action carries regulatory and financial weight, is where most programs slow down or stall entirely. That's precisely why getting it right creates separation that competitors can't easily copy.

So the competitive question isn't "Are you using AI?" Almost everyone is. The real question is "Can you scale it into your core operations safely, and prove it?" Far fewer carriers can answer yes.

The governance gap

The distance between a successful pilot and safe operation at scale is wider than most teams expect. And the gap isn't really about the technology. It's about governance.

Most AI pilots don't need much governance because the AI can't cause much harm. It recommends, and a human decides. Once an agent can take action inside a regulated system, the picture changes completely. Suddenly you need clear answers to questions that never came up in the sandbox.

Four issues tend to surface at the same time.

01 Governance was never designed in

Because early experiments lived at the edges of the business, they were never built to answer questions like: Who authorized this action? Under what policy? What were the boundaries? Many carriers discover, at exactly the wrong moment, that their AI initiatives can't produce those answers. The governance model has to exist before the agent acts, not after something goes wrong.

02 Regulatory scrutiny is intensifying

Insurance is among the most heavily regulated industries anywhere. Carriers answer to state and national regulators, solvency requirements, data protection rules, and fair-treatment standards. Regulators are now looking closely at automated decision-making, including explainability, bias, and consumer fairness. An AI action that can't be explained or defended isn't an efficiency. It's exposure.

03 Compliance obligations don't pause for AI

Core insurance processes carry documentation and control obligations. Claims decisions must be justifiable. Underwriting must apply consistent, defensible criteria. Rate changes must follow approved rules. When an agent participates in any of these processes, its actions become part of the compliance record. If the system can't produce clean evidence of what it did and why, it introduces risk a compliance team simply can't accept.

04 Core systems have little tolerance for surprises

Core insurance systems are often decades in the making. They span mainframe, distributed, and cloud environments, and they run the workflows a carrier can't afford to break. Nightly batch cycles, policy renewals, and settlement runs depend on precise, predictable behavior. An agent that improvises in that setting is a liability, not a helper.

Put these four together and the pattern is clear. The governance gap, not the intelligence gap, is what keeps insurers stuck between promising pilots and real operational value. The carriers who close that gap first will be the ones who scale.

What scaling AI safely actually requires

If governance is the real barrier, then closing it deliberately becomes the path to advantage. Scaling agentic AI safely across insurance operations rests on five requirements. None of them are optional in a regulated environment.

Policy-based control

Actions should be constrained by explicit policy, expressed in a form the system enforces rather than merely documents. Not every action is equally permissible, and the AI agent needs to understand and respect those limits before it does anything. Policy-as-code turns written rules into rules the system actually follows.

Identity-aware execution

The system should always know who initiated or approved an action. Access should be tied to role and authority, so an AI agent never operates beyond the permissions of the person behind the request. An action taken on behalf of a junior analyst should carry that analyst's limits, not an administrator's.

Auditability

Every meaningful action should be traceable after the fact. Leaders, auditors, and regulators need to understand what happened, when it happened, why it happened, and who approved it. In insurance, this is often the difference between a defensible decision and a regulatory finding.

Human oversight

Some actions require review or explicit approval before they run. Human involvement should be built into the operating model, not treated as an exception. When risk crosses a defined threshold, the system should stop and escalate rather than proceed on its own.

Deterministic execution

Critical actions must follow defined rules and validations rather than improvisation. In plain terms, the system follows established steps instead of inventing new ones where precision matters. For a settlement run or a rate application, predictability is the whole point.

What makes these requirements a competitive lever rather than a constraint is that they're what allow you to expand. Once governance is engineered in, you can add use cases with confidence instead of relitigating trust every time. That's how a handful of AI pilots becomes a scaled program.

Where EVA fits

Rocket EVA is a governed agentic AI platform built for enterprise operations, including the complex environments common in insurance. It's designed for settings where control, compliance, and reliability matter as much as speed.

In plain terms, EVA helps operations teams diagnose issues, trace root causes, and support action across systems that span mainframe, distributed, and cloud environments. It works inside existing security and operational frameworks rather than around them. It sits alongside the tools a carrier already runs and relies on existing identity, policy, and approval layers to do its work.

It helps to be clear about what EVA is and isn't. It's a diagnostic and action-support layer that operates under governance. It isn't an autonomous decision-maker that replaces underwriters, adjusters, or system administrators. It doesn't move your data out of the systems where it already lives, and it doesn't act outside the boundaries your controls define. For a systems team, the practical change is that time once spent hunting through logs and correlating symptoms across siloed tools shifts toward reviewing clear, governed findings.

EVA's design maps directly to the five requirements above. Governance is built in through policy-as-code and human approval where it's needed. Execution follows controlled, validated paths rather than improvisation. A context graph links applications, jobs, transactions, infrastructure, and services so the system can trace an issue from symptom toward root cause. And an open tool layer lets it integrate with the identity, service management, observability, and automation systems a carrier already depends on. Those design choices are what make scaling feasible, because each new use case inherits the same guardrails rather than starting from scratch.

Insurance use cases at scale

The value becomes concrete in the day-to-day work of running insurance systems, and it compounds as you extend it across more of the operation.

Claims processing at peak volume

When a claims platform slows or throws errors during a high-volume period, adjusters stall and customers wait. Tracing the cause often means jumping across databases, application servers, integration layers, and sometimes the mainframe behind them. EVA correlates signals across those layers to move from symptom toward likely root cause, which can shorten investigation time. High-risk remediation still routes to an authorized human, and every step stays logged for later review. Governance matters here because a claims system is a system of record, and unexplained changes to it create compliance exposure. Scaled across multiple claims lines, that same governed pattern holds without adding review overhead each time.

Underwriting system diagnostics

Underwriting workbenches pull data from many sources, and a single broken feed can quietly degrade decisions across a book of business. EVA can help trace where a data pipeline failed or a service dependency broke, so teams find the problem before it spreads. Because access is identity-aware, the platform only surfaces and acts within the permissions of the person driving the work, which keeps sensitive rating and applicant data governed by existing controls.

Nightly batch operations

Batch is the heartbeat of insurance operations. Policy renewals, premium calculations, commission runs, and regulatory reporting often depend on jobs that run overnight and must finish before morning. A single failed job can cascade. EVA can identify recurring failure patterns, classify exceptions, and prioritize the ones that threaten the critical path, reducing the manual triage a team faces at 3 a.m. Deterministic execution matters here, because the system follows defined, validated steps rather than improvising against production jobs.

Cross-system incident diagnostics

Major incidents rarely respect system boundaries. A slowdown may start in the network, surface in a database, and manifest in a customer-facing application. EVA traces issues across network, database, application, and mainframe layers to shorten root cause analysis and improve response. Throughout, human oversight and audit trails keep the response defensible, which is exactly what a regulator or internal auditor will later want to see.

Across all four, the same principle holds. EVA accelerates the diagnosis and supports the action, but it does so inside a governed model where authority, approval, and evidence are never optional. That consistency is what lets a carrier extend from one use case to many without multiplying risk.

Why governance will separate winners from laggards

Step back and the competitive picture comes into focus. Nearly every insurer can run a pilot. Far fewer can put an agent into a claims platform or a batch schedule and prove, months later, that every action it took stayed within policy and remained traceable. Fewer still can do that across dozens of use cases at once.

That's the dividing line. Insurers who stall tend to treat governance as something to sort out after the technology works. Insurers who scale treat governance as the thing that lets the technology work at all. They start with a focused, lower-risk use case, keep existing controls in place, and expand as confidence grows. They don't ask their AI to do everything. They ask it to do specific things safely, prove it, and then do more.

The payoff for getting this right is meaningful and cumulative. Faster incident resolution and lower mean time to resolution. Less manual triage consuming scarce expert time. Better use of experienced staff as institutional knowledge is preserved rather than lost. And a stronger position with regulators, because automation backed by clear evidence is far easier to defend than automation that can't explain itself.

None of these advantages come from moving fastest. They come from being able to move safely, repeatedly, and at scale. That's a capability competitors can't replicate overnight, which is exactly what makes it a durable advantage rather than a fleeting one.

Trust is the long-term differentiator

The next phase of AI in insurance won't be won on speed or novelty. Those advantages are real, but they're easy to copy and quick to fade. The advantage that lasts belongs to carriers who can scale AI into their core systems and prove, at any moment, that it's operating within the rules.

Trust will define the winners of agentic AI in insurance. Not trust as a marketing message, but trust as an engineered property of the system: policy that the platform enforces, identity that constrains what it can do, evidence it produces without being asked, and human oversight where the stakes are high. Carriers who build these guardrails first will scale with confidence while others are still trying to figure out whether their pilots can survive contact with production.

Governed agentic AI at scale isn't the cautious path. It's the competitive one. The insurers who understand that will be the ones still leading when the rest of the industry finally has to answer the hard questions.

Learn more about **Rocket EVA** **or** **Speak to an expert today**

About Rocket Software

Rocket Software is a global technology leader in modernization and a partner of choice that empowers the world's leading businesses on their modernization journeys, spanning core systems to the cloud. Trusted by over 12,500 customers and 750 partners, and with more than 3,200 global employees, Rocket Software enables customers to maximize their data, applications, and infrastructure to deliver critical services that power our modern world. Rocket Software is a privately held U.S. corporation headquartered in the Boston area with centers of excellence strategically located throughout North America, Europe, Asia and Australia. Rocket Software is a portfolio company of Bain Capital Private Equity.

Rocket® software

Modernization. **Without Disruption.**™

Visit RocketSoftware.com

MAR-18236_WP_EVAInsuranceThoughtLeadership_V2