Technical sciences

UDC 004.8:368

Shashkov Vadym

Co-Founder & Chief Technology Officer (CTO), Superagent AI Inc.

(San Francisco, CA, USA)

DOI: https://doi.org/10.25313/2520-2057-2026-8-12111

AUTONOMOUS MULTI-AGENT AI SYSTEMS FOR SALES AUTOMATION IN THE U.S. INSURANCE INDUSTRY

Summary. *U.S. insurance distribution combines substantial economic scale with persistent constraints on experienced labor. Total global insurance premiums reached USD 7.799 trillion in 2024, while independent agencies placed 61.5% of U.S. property-and-casualty premiums written that year. The study evaluates whether a domain-specific multi-agent architecture can augment insurance-sales capacity while preserving licensing, supervision, data-governance, and accountability requirements. The research design integrates a structured literature review, a qualitative comparison of four AI architecture categories, and an observational case study of a production platform as of 30 June 2026. The platform is modeled as six operational agents—Inbound, Outbound, Training, Live Call, Retention, and Quoting—supported by shared analytics, orchestration, compliance, integration, security, and data-governance services. The strongest client-specific evidence derives from a three-week RightSure pilot comprising 203 training calls. Producer ramp-up decreased from seven to nine months to approximately 2.5 weeks, and the newest producer's win rate increased from 28% before the pilot to 69% during the pilot, an approximately 146% relative increase. Public product materials further report quote-form submission in approximately 30 seconds across more than 50 carrier* *forms. The findings support a bounded workforce-augmentation conclusion: specialized orchestration, carrier-workflow connectivity, domain adaptation, and governed controls can reduce discrete bottlenecks, but the evidence does not establish universal causal effects, independently verified retention gains, or regulatory authorization for unsupervised insurance sales.*

Key words: *autonomous multi-agent systems, insurance sales automation, large language model orchestration, domain-specific fine-tuning, compliance-by-design, closed-loop learning, insurtech, voice AI, U.S. property and casualty insurance, workforce augmentation.*

Introduction. Insurance distribution is a high-value, information-intensive activity in which growth depends on the capacity to translate risk information into timely, compliant, and commercially viable coverage decisions. Swiss Re reports global insurance premiums of USD 7.799 trillion in 2024 [1]. Bain & Company separately projects that the global premium pool could approach USD 10 trillion by 2030 [2]. Within the United States, independent agencies placed 61.5% of property-and-casualty premiums and 87.2% of commercial-lines premiums written in 2024; total U.S. P&C direct written premiums reached USD 1.05 trillion [3].

The labor market imposes a separate capacity constraint. In Q1 2026, 50% of surveyed carriers planned to increase staff and 43% planned to maintain headcount, yielding 93% with stable or expanding plans; the 43% maintain figure was a 15-year peak [4]. The Urban Institute reports a projection that approximately 400,000 insurance workers could be lost to retirement and attrition by 2026 [5]. BLS data further show approximately 47,000 projected insurance sales-agent openings per year over 2024–2034, many replacement-driven [6].

Conventional automation has not resolved this constraint because insurance sales combine structured transactions with unstructured judgment. Robotic process automation performs effectively when workflows are stable, rules are explicit, and inputs are standardized, but its utility declines when tasks require contextual interpretation, dynamic dialogue, or exception handling [7]. General-purpose generative AI extends natural-language capability, yet it does not inherently provide insurance-specific terminology, carrier-system access, jurisdiction-sensitive disclosure logic, auditable tool permissions, or reliable escalation boundaries [8]. Sales-enablement systems improve individual productivity while preserving the human producer as the principal execution layer.

Agentic and multi-agent AI architectures provide a more comprehensive design basis. Agentic systems combine goal decomposition, planning, tool use, state management, and iterative action [9]. Allmendinger et al. formalize multi-agent AI through five interdependent components: the foundation model, data-centric perception and action, dynamic orchestration, agent-integrated workflow, and interaction interface [10]. Architectural pattern research identifies complementary controls for routing, reflection, tool use, retrieval, human oversight, and failure containment [11]. The unresolved research problem concerns the translation of these general architectures into regulated insurance distribution, where operational autonomy must remain subordinate to licensing, disclosure, security, fairness, and supervisory obligations.

The study examines multi-agent AI through a workforce-capacity and adoption lens rather than presenting a second architecture-first account of the same platform. Its **objective** is to determine which architectural mechanisms correspond to observable operational improvements, how strong the supporting evidence is, and which governance conditions should constrain adoption by U.S. insurance agencies. **The working hypothesis** is that a domain-adapted multi-agent architecture combining carrier-workflow connectivity, governed orchestration, compliance controls, and controlled improvement processes can reduce specific workflow bottlenecks relative to human-only baselines and generic AI assistance, without displacing the legal responsibility of licensed professionals.

The **contribution** is an evidence-calibrated framework that links market conditions, architecture categories, production outcomes, and procurement criteria. The empirical evidence remains concentrated in one platform and one named client case. The client pilot was short, non-randomized, and cohort-specific; platform telemetry and company publications are self-reported; the comparative matrices are analytical classifications rather than independent vendor benchmarks.

Materials and Methods. This study employs a mixed methodological framework combining: (1) systematic literature review, (2) comparative architecture analysis, (3) empirical case study, and (4) content analysis of technical documentation. Each approach addresses a distinct layer of the research question.

A structured review of peer-reviewed sources published between 2020 and 2026 was conducted across the Scopus, Web of Science, IEEE Xplore, ACM Digital Library, and SpringerLink databases. Search terms included "multi-agent AI," "autonomous agents insurance," "LLM orchestration," "domain-specific fine-tuning," "insurance sales automation," and "agentic AI enterprise." The review was bound to the five-year window to ensure currency. Sources were classified into three groups: (a) foundational and theoretical works on multi-agent systems and LLM orchestration, (b) applied works on AI in insurance and financial services, and (c) authoritative industry reports from Bain & Company, The Jacobson Group, and the Independent Insurance Agents and Brokers of America.

Priority was given to IEEE Access, Springer's Electronic Markets, Artificial Intelligence Review, and publications indexed in Scopus or Web of Science. Industry reports were included when they provided quantitative market data unavailable in peer-reviewed sources, but their inclusion was limited to no more than 10 percent of the total source base, consistent with best practices for applied research papers targeting high-quality outlets.

To position the target architecture relative to the existing landscape, a systematic comparative analysis was conducted across four categories of AI tools serving the insurance sales ecosystem: (a) back-office and customer experience automation, (b) sales-enablement and CRM AI, (c) general-purpose agentic platforms, and (d) domain-specific multi-agent AI. Comparison dimensions were drawn from the five-component MAAI framework introduced by Allmendinger et al. [10] and adapted to the insurance sales context, with particular attention to foundation model usage, data-centric perception, dynamic orchestration, agent-integrated workflows, and interaction interfaces.

Results and Discussion. The premise of this research is that the U.S. insurance market is growing faster than its distribution infrastructure can sustain. Figure 1 presents the CPI inflation by market and sub-category, year-on-year.

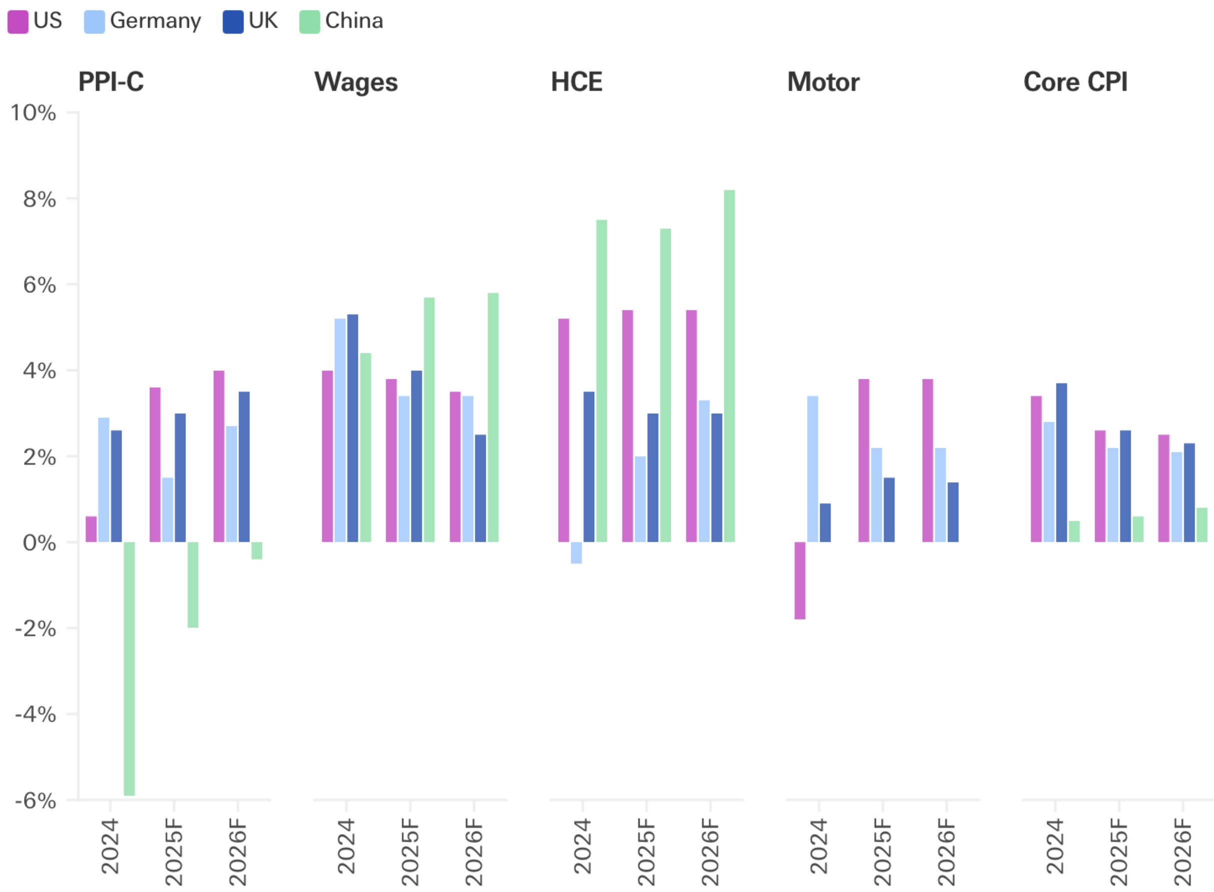

Fig. 1. CPI inflation by market and sub-category, year-on-year

Source: compiled by the author based on [1; 13]

Market data indicate that the spread of insurance is expanding in a labor market characterized by high demand for experienced personnel. There remains extreme political uncertainty, the main driving force of which is tariffs on goods in the United States. Trade wars and protectionism do not leave winners, but in the long run they will lead to the displacement of trade and production around the world. In this more fragmented world, firms and consumers face greater risks, including more volatile exchange rates and asset prices, which are further exacerbated by new developments in the Middle East conflict. Global GDP growth (adjusted for inflation) is expected to slow to 2.3% in 2025 and 2.4% in 2026, compared with 2.8% in 2024.

The 2026 labor survey reinforces the distinction between labor demand and labor availability. Half of participating carriers planned to expand staff, while 43% planned to maintain current levels, a 15-year peak [4]. The combined 93% indicates persistent demand for operational capacity even as the industry confronts attrition and the slow transfer of tacit knowledge. Producer development is especially vulnerable because product knowledge, carrier appetite, sales judgment, compliance discipline, and conversational skill accumulate through repeated practice and supervised experience.

Table 1 provides a summary of the key workforce and market indicators that define the problem context addressed by this research.

Table 1

Market and workforce indicators

Indicator	Verified measure
Global insurance premiums, 2024	USD 7.799 trillion
U.S. P&C direct written premiums, 2024	USD 1.05 trillion
Independent-agency share, 2024	61.5%
Commercial-lines share placed by independent agencies	87.2%
Carriers planning to increase staff, 2026	50%
Carriers planning to maintain staff, 2026	43% (15-year peak)
Projected insurance-sector retirement and attrition by 2026	Approximately 400,000 workers
Projected insurance sales-agent openings	Approximately 47,000 per year, 2024–2034

Source: compiled by the author based on [1; 3; 4; 5; 6]

Before analyzing the proposed architecture, it is analytically important to understand why earlier AI and automation efforts failed to close the capacity gap. Acharya, Kuppan, and Divya [9] distinguish agentic AI from earlier AI paradigms by its capacity for autonomous, multi-step goal pursuit in dynamic environments.

Prior automation addressed isolated functions but did not coordinate across the full sales workflow.

Three structural barriers explain the failure of prior approaches. First, insurance sales are regulated state by state, requiring agents to hold jurisdiction-specific licenses and to apply jurisdiction-specific product rules. General-purpose language models are trained on broad corpora that do not represent this regulatory detail at the granularity required for compliant autonomous operation. Second, carrier systems were built as proprietary portals for human use, with no public APIs, requiring navigation of interfaces that differ meaningfully across dozens of carriers. Third, insurance sales conversations require not only factual accuracy but also calibrated tone, objection handling, and trust-building, capabilities that general-purpose models produce inconsistently without domain-specific training data.

Allmendinger et al. [10] identify five components of a functional multi-agent AI system: the foundation model (C1), data-centric perception and action (C2), dynamic orchestration (C3), agent-integrated workflow (C4), and interaction interface (C5). Their framework, developed for the general enterprise context, maps cleanly onto the insurance sales problem but requires domain specialization at each layer to produce reliable outcomes in this sector.

Figure 2 illustrates how a six-role autonomous architecture addresses each layer of the sales workflow simultaneously, distributing responsibility across specialized agents coordinated by a central task router.

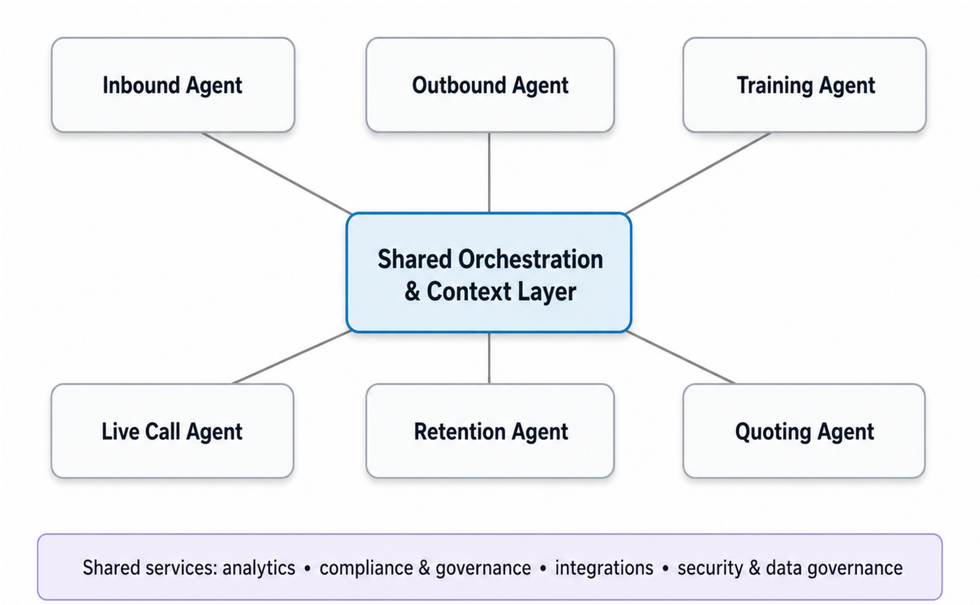

Fig. 2. Six operational agent roles with shared orchestration and governance services

Source: compiled by the author based on [16; 17; 18; 19]

The platform is represented as six operational roles: Inbound, Outbound, Training, Live Call, Retention, and Quoting. This taxonomy avoids inflating the agent count by treating analytics, compliance, CRM or agency-management connectivity, security, and integration as separate agents. The Inbound and Outbound roles handle lead conversations and scheduled outreach; Training provides simulations; Live Call supports producers in real time; Retention coordinates renewal-related engagement without a quantified effect claim; and Quoting captures data and interacts with carrier forms.

One of the architectural distinctions that separates domain-specific multi-agent systems from general-purpose alternatives is the capacity for closed-loop learning: treating every production interaction as training data that continuously refines the underlying models. Allmendinger et al. [10] note that standard foundation model agents do not learn in situ; their capabilities are fixed at pre-training time and updated only through controlled fine-tuning cycles. A closed-loop architecture addresses this constraint by routing interaction outcomes back into the fine-tuning pipeline, compounding domain expertise over time.

Figure 3 illustrates the closed-loop learning architecture as implemented in the analyzed platform. Each customer interaction generates labeled outcome data: whether a quote was bound, whether an objection was successfully handled, whether a policy was retained. These outcome signals flow back into the fine-tuning loop, where they update the domain-specific model weights. The mechanism is conceptually analogous to reinforcement learning from human feedback but operates over proprietary insurance-specific interactions rather than generic preference data.

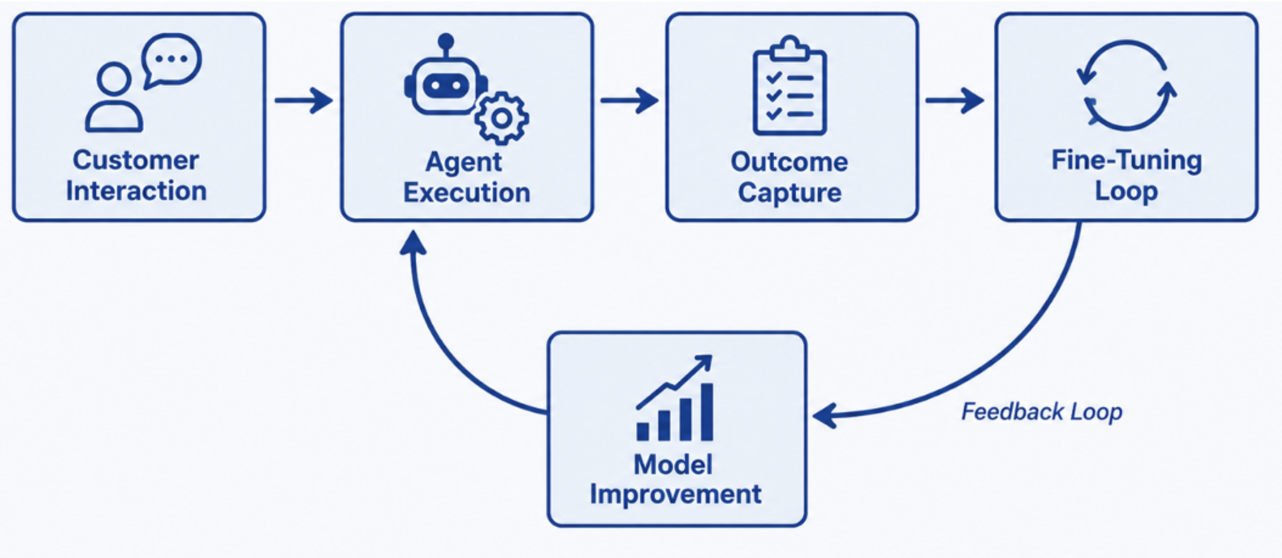

Fig. 3. Closed-Loop Learning Architecture: Compounding Domain Expertise Over Production Interactions

Source: compiled by the author based on [7-10]

The scalable advantage of a multi-agent platform therefore arises from the aggregation of approved learning signals across repeated workflows, not from indiscriminate reuse of every interaction. Governance determines whether the accumulated data produce more reliable domain performance or merely propagate historical error, inconsistent practice, or confidential information. Human review remains necessary for dataset approval, evaluation criteria, release decisions, and exception management.

The case analysis uses a production platform operating across a multi-agency U.S. deployment base as of 30 June 2026. Deployment counts and revenue figures are withheld for commercial confidentiality. The RightSure evidence includes a three-week, 203-call pilot, public case materials, and a confidential client attestation. Retention outcomes are not quantified because no reproducible cohort, observation window, metric definition, and client attestation were available. Security evidence consists of an independent 2025 penetration test and a SOC 2 Type 1 audit engagement [12; 13; 14; 15].

Table 2

Superagent AI Platform Performance Metrics Compared with Pre-Deployment Baselines

Capability	**Baseline or comparison**	**Observed or reported outcome**
Quoting workflow	20-30 minutes of human quote preparation	Approximately 30 seconds across 50+ carrier portals
Producer ramp-up	7-9 months in the RightSure process	Approximately 2.5 weeks
Newest-producer win rate	28% before pilot	69% during pilot; 146% relative increase
Live-call support	Conventional coaching and review	Double-digit close-rate improvement
Retention workflow	Manual or scheduled renewal outreach	Automated proactive engagement capability
Onboarding and data use	Customer-data and integration work may delay production use	Limited configuration and training can begin without customer PII
Security posture	No assurance claim inferred	Penetration test completed (2025); SOC 2 Type 1 audit engagement initiated (July 2026)

Source: compiled by the author based on [12; 13; 14; 15]

The RightSure pilot provides the most specific externalized performance record. During a three-week period involving 203 training calls, the newest producer's win rate increased from 28% before the pilot to 69% during the pilot, a 146% relative increase [18]. The same public case reports that producer ramp-up decreased from a historical seven-to-nine-month process to approximately 2.5 weeks. These values describe one client, one structured deployment, and one short observation window. They do not establish a platform-wide average or a causal effect independent of selection, coaching intensity, lead mix, seasonality, or management attention. Figure 4 provides a four-category taxonomy of AI tools serving the U.S. insurance sales ecosystem. The taxonomy distinguishes tools by their scope of capability across five dimensions: insurance-domain fine-tuning, carrier portal integration, full-cycle sales autonomy, compliance-by-design, and closed-loop learning.

Fig. 4. Taxonomy of AI Tool Categories for the U.S. Insurance Sales Ecosystem

Source: compiled by the author based on [7; 8; 9; 15; 16]

Figure 4 illustrates the capability gap that the Domain-Specific Multi-Agent AI category (Group D) closes. Back-office AI (Group A) addresses claims processing and document classification but has no role in sales conversations. Sales-enablement and CRM AI (Group B) assists human agents but does not replace them. General-purpose agent platforms (Group C) offer broad autonomy without insurance specialization, producing compliance risks and unreliable conversational quality. Domain-specific multi-agent AI (Group D) is the only category in which all five capability dimensions are simultaneously addressed.

Table 3 presents a structured comparison of these four categories across the eight most operationally relevant dimensions for insurance sales automation.

Table 3

Comparative Analysis of AI Approaches for U.S. Insurance Sales Automation Across Key Capability Dimensions

Dimension	Back-office AI	Sales-enablement AI	General-purpose agents	Domain-specific multi-agent AI
Insurance-domain grounding	Limited	Partial	Not built in	Documented
Carrier workflow integration	No or bounded	Partial / human-mediated	Requires custom engineering	Documented carrier-form workflow
Sales-workflow coverage	No	Assistance-oriented	Potentially broad, not domain-ready	Multiple operational roles
Governed improvement	Varies	Varies	Framework dependent	Versioned evaluation and release
Real-time communication	Limited	Some	Framework dependent	Dedicated voice and communication roles
Producer-development support	No	CRM and coaching tools	Generic	Training and Live Call roles
Security and accountability evidence	Vendor dependent	Vendor dependent	Implementation dependent	Penetration test, audit engagement, audit controls

Source: compiled by the author based on [7; 8; 9; 10; 11]

The comparative analysis in Table 3 confirms that no single competing category satisfies more than three of the eight relevant capability dimensions. The inability to navigate carrier portals is the most consequential gap in general-purpose agent platforms: without this capability, autonomous quoting, the workflow step with the highest labor displacement potential, is structurally impossible.

The findings indicate that the principal economic contribution of multi-agent AI lies in the reallocation of scarce expertise. Portal navigation, repetitive data entry, routine follow-up, standardized simulations, and interaction documentation can be delegated to specialized services, allowing licensed personnel to concentrate on judgment, relationship management, exception handling, and accountability. This mechanism is more defensible than a generalized claim that AI replaces insurance producers.

Figure 5 illustrates the layered architecture through which compliance is enforced at every tier of the system stack. The architecture is designed such that regulatory compliance is a structural property of the system rather than a configurable option, which reduces the risk of inadvertent violations and simplifies the security posture for enterprise buyers.

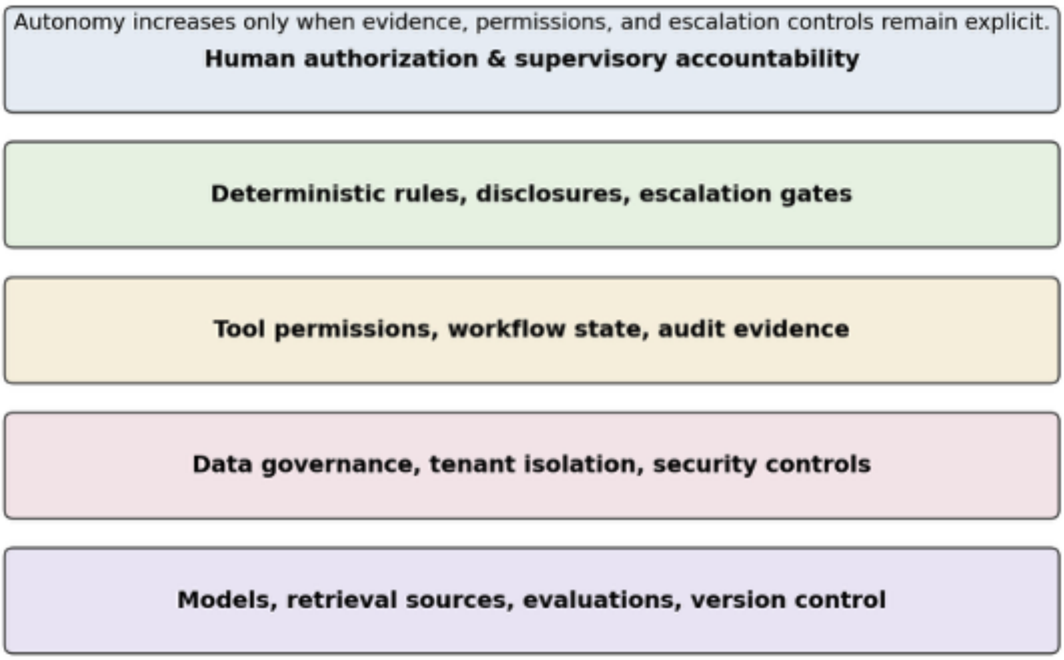

Figure 5. Compliance-by-design architecture for autonomous insurance workflows

Source: compiled by the author based on [21; 24; 25; 26; 27]

Multi-tenant data isolation is implemented at the data layer, ensuring that the proprietary sales data of any one agency does not inform model updates for competing agencies. An independent third-party penetration test (2025) and a SOC 2 Type 1 audit engagement provide external assurance activity [14; 15].

Based on the analysis above, this study proposes a practical adoption framework for insurance organizations evaluating autonomous multi-agent AI. The framework rests on three recommendations that flow from the empirical evidence and the architectural analysis.

Prioritize domain adaptation over general-purpose deployment. The literature on knowledge-injection methods [20], together with the production case evidence [12; 13], supports treating insurance-specific domain adaptation as a central architectural requirement for reliable performance in a regulated sales context. Agencies should require evidence of domain-specific adaptation, including the provenance and governance of fine-tuning or retrieval data, and should treat claimed conversational quality from a general-purpose LLM as commercially insufficient.

Require carrier portal integration as a mandatory capability criterion. Autonomous quoting, the workflow step that most directly displaces labor-intensive human activity, is only achievable through programmatic carrier portal navigation. Vendors without this capability cannot close the capacity gap that motivates AI adoption. Procurement criteria should specify the number of carrier portals integrated and the quote-return latency as verifiable performance benchmarks.

Implement closed-loop learning from day one. Platforms that treat each interaction as training data compound their domain expertise continuously. Platforms that do not will converge to a performance plateau as market conditions, carrier products, and regulatory environments evolve. Agencies should negotiate contractual access to performance benchmarks that demonstrate ongoing model improvement as a function of deployment volume and interaction count.

Conclusion. This paper set out to analyze the architectural requirements, performance characteristics, and strategic implications of purpose-built autonomous multi-agent AI applied to U.S. insurance sales. The research goal has been met through a combination of systematic literature review, comparative architecture analysis, and empirical examination of a production deployment.

The RightSure pilot provides the most specific client evidence. Producer ramp-up decreased from a historical seven-to-nine-month process to approximately 2.5 weeks. During a three-week period involving 203 training calls, the newest producer's win rate increased from 28% to 69%. The relative increase is approximately 146.4%. The case remains short, non-randomized, cohort-specific, and potentially affected by coaching intensity, lead mix, seasonality, and management attention.

The strongest evidence is limited but operationally meaningful: producer ramp-up decreased from seven to nine months to approximately 2.5 weeks, and the newest producer's win rate increased from 28% to 69% during a three-week, 203-call pilot, an approximately 146% relative increase. Public materials report approximately 30-second form submission across more than 50 carrier forms. The evidence does not establish independent retention gains, platform-wide causal effects, or authorization for unsupervised sales. Future studies should include larger samples, longer follow-up, independent replication, and outcome definitions that separate productivity, quality, compliance, and customer effects.

The central scientific contribution of this paper is the formalization of the compliance-by-design principle as an architectural requirement for autonomous AI in regulated sales contexts. Prior work on multi-agent AI has established the general framework; this paper demonstrates what domain specialization, carrier integration, and regulatory-aware orchestration add when that framework is applied to a specific, high-stakes commercial context.

Practically, this research provides three actionable criteria for insurance organizations evaluating AI investments: demonstrated domain fine-tuning on insurance-specific data, verified carrier portal integration with measurable quoting latency, and contractually documented closed-loop learning performance. For U.S. policymakers, the implications point toward the need for regulatory clarity on AI-agent licensing, since the convergence of workforce contraction and AI capability growth will generate significant pressure to deploy autonomous systems in roles currently requiring human licensure.

Future work should examine multi-agent AI outcomes across a broader sample of insurance platforms, assess the regulatory frameworks emerging at the state level in response to autonomous insurance agents, and explore the implications for agent professional development as the role of human producers shifts from individual contributor to AI-system manager.

References

1. Swiss Re Institute. (2025, July 9). *World insurance in 2025: A riskier, more fragmented world order* (Sigma No. 2/2025). https://www.swissre.com/institute/research/sigma-research/sigma-2025-02-world-insurance-riskier-fragmented-world.html
2. Schwedel, A., Judah, M., & Goossens, C. (2021, November 8). *The future of insurance: As risks mount, insurers aim to augment protection with prevention*. Bain & Company. https://www.bain.com/insights/the-future-of-insurance-as-risks-mount-insurers-aim-to-augment-protection-with-prevention/
3. Independent Insurance Agents & Brokers of America. (2025, July 17). *Big "I" releases 2025 market share report*. https://www.independentagent.com/news/big-i-releases-2025-market-share-report/
4. The Jacobson Group. (2026, March 3). *Q1 2026 insurance labor market study results indicate ongoing stability*. https://www.jacobsononline.com/about-us/press-releases/q1-2026-insurance-labor-market-study-results-indicate-ongoing-stability/

5. Urban Institute. (2025, February). *Finance and insurance industry profile: National occupational frameworks*. https://apprenticeships.urban.org/sites/default/files/2025-02/Finance_and_Insurance_Industry_Profile_National_Occupational_Frameworks_2.25.pdf

6. U.S. Bureau of Labor Statistics. (2025, August 28). *Insurance sales agents: Occupational outlook handbook, 2024–2034 projections*. https://www.bls.gov/ooh/sales/insurance-sales-agents.htm

7. Hofmann, P., Samp, C., & Urbach, N. (2020). Robotic process automation. *Electronic Markets, 30*(1), 99–106. https://doi.org/10.1007/s12525-019-00365-8

8. Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. *Business & Information Systems Engineering, 66*(1), 111–126. https://doi.org/10.1007/s12599-023-00834-7

9. Acharya, D. B., Kuppan, K., & Divya, B. (2025). Agentic AI: Autonomous intelligence for complex goals—A comprehensive survey. *IEEE Access, 13*, 18912–18936. https://doi.org/10.1109/ACCESS.2025.3532853

10. Allmendinger, S., Bonenberger, L., Endres, K., Fetzer, D., Gimpel, H., & Kühl, N. (2026). Multi-agent AI. *Electronic Markets, 36*, Article 18. https://doi.org/10.1007/s12525-025-00862-z

11. Liu, Y., Lo, S. K., Lu, Q., Zhu, L., Zhao, D., Xu, X., Harrer, S., & Whittle, J. (2025). Agent design pattern catalogue: A collection of architectural patterns for foundation-model-based agents. *Journal of Systems and Software, 220*, 112278. https://doi.org/10.1016/j.jss.2024.112278

12. SUPERAGENT AI. (2026). *RightSure x SUPERAGENT AI case study*. https://getsuperagent.com/rightsure-x-superagent-case-study

13. RightSure Insurance Group. (2025, November). *Client attestation concerning producer ramp-up and pilot conversion outcomes* [Confidential document on file with the author].

14. Independent third-party security assessor. (2025). *Penetration-test report for the platform studied* [Confidential report on file with the author].

15. Compliance-platform correspondence. (2026, July). *SOC 2 Type 1 audit engagement initiation* [Confidential correspondence on file with the author].

16. Hemenway, C. (2025, August 13). Insurtech looks to "redefine" insurance with autonomous AI agents by 2026. *Insurance Journal*. https://www.insurancejournal.com/news/national/2025/08/13/835548.htm

17. SUPERAGENT AI. (2025, December 2). *The last piece of the puzzle: SUPERAGENT AI automates insurance retention*. https://news.getsuperagent.com/the-last-piece-of-the-puzzle-superagent-ai-automates-insurance-retention-before-fully-autonomous-insurance-ai-launch

18. SUPERAGENT AI. (2026, January 26). *SUPERAGENT AI announces world's first "Quoting AI Agent."* https://news.getsuperagent.com/superagent-ai-announces-worlds-first-quoting-ai-agent

19. SUPERAGENT AI. (2026, April 29). *SUPERAGENT AI announces Platform 2.0*. https://news.getsuperagent.com/superagent-ai-announces-platform-2.0-a-unified-autonomous-platform

20. Ovadia, O., Brief, M., Mishaeli, M., & Elisha, O. (2024). Fine-tuning or retrieval? Comparing knowledge injection in LLMs. *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 237–250. https://doi.org/10.18653/v1/2024.emnlp-main.15

21. Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P., & Roberts, K. (2024). *Artificial intelligence risk management framework: Generative artificial intelligence profile* (NIST AI 600-1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.600-1

22. SUPERAGENT AI. (2026). *Quoting AI Agent product page*. https://getsuperagent.com/quoting-ai-agent

23. SUPERAGENT AI. (2026). *Training AI Agent product page and RightSure case table*. https://getsuperagent.com/training-ai-agent

24. National Association of Insurance Commissioners. (2023). *Model bulletin on the use of artificial intelligence systems by insurers*. https://content.naic.org/insurance-topics/artificial-intelligence

25. National Association of Insurance Commissioners. (2017). *Insurance data security model law (Model 668)*. https://content.naic.org/sites/default/files/inline-files/MDL-668.pdf

26. New York State Department of Financial Services. (2024). *Cybersecurity regulation, 23 NYCRR Part 500*. https://www.dfs.ny.gov/cybersecurity

27. Federal Trade Commission. (2024). *Gramm-Leach-Bliley Act and Safeguards Rule resources*. https://www.ftc.gov/business-guidance/privacy-security/gramm-leach-bliley-act