INSIGHT PAPER

Ceded Reinsurance in the Age of Autonomous Insurance

What changes, what accelerates, and what your operation needs before you deploy agentic AI.

Susanne Schmitz, Principal Consultant, Reinsurance, Sapiens

sapiens

Executive Summary

If you run ceded reinsurance at a P&C insurer, you know the friction. Recoveries turn up late or not at all. Month-end close takes days longer than it should. Schedule F gets built by hand. And at renewal, the quality of your data shapes the conversation before you say a word.

Agentic AI changes that model by acting on your workflows directly, working toward the objectives you give it. The insurers seeing real results led with operating model work, not the technology. This paper covers both sides: what agentic AI delivers for ceded reinsurance, and what your team needs in place before it delivers consistently.

The reality your team is living with

A Deloitte survey on ceded reinsurance found 71% of participating insurers say late or incomplete data drives manual workarounds, and two-thirds of ceded data processing still involves manual handling. The cost lands in two places.

Recoverable leakage: Claim data that doesn't attach to the right treaty means recoveries fall through systematically. Audits find shortfalls of $1.3 million to over $10 million per organization – amounts that were always owed, just never collected.

The reconciliation "tax": Skilled accountants rebuilding bordereaux from scratch each period, resolving discrepancies that clean data architecture would have prevented. The annual cost runs $290,000 to $1.6 million depending on portfolio size.

Both point to a broader issue: if reinsurance costs have risen even as loss ratios stayed stable, your treaty pricing may embed a data quality premium that reinsurers are charging but not naming.

> *Without reliable data, AI mainly creates additional complexity, costs and frustration.*
>
> – **Andreas Berger, CEO of Swiss Re, April 2026**

How AI amplifies what is already true

Conventional AI tools made existing operating model conditions more consequential. Where treaty, claims, and financial data were already connected, AI produced better decisions faster. Where they were fragmented, it produced confident wrong answers at higher volume. The operating model was the variable. The AI was the amplifier.

Agentic AI operates on the same principle, with higher stakes. It plans, decides, and acts autonomously across multi-step workflows. That capability is already running at your competitors. Insurance AI deployments grew 87% year over year according to Evident, with agentic AI at 21% of publicized deployments in Q4 2025 and commercial P&C insurers reporting loss ratio improvements of 3 to 5 percentage points from systems already running. Reported outcomes, not forecasts.

> *Conventional AI works on your operating model. Agentic AI works inside it.*

What agentic AI actually changes for you

Generative AI waits to be asked. Agentic AI acts. For ceded reinsurance – a process-execution function, not a question-answering one – three capabilities make the practical difference.

Understand

Agents read slip letters, cover notes, bordereaux in any format, and extract structured, treaty-aware data without anyone re-keying it.

Reason and decide

Agents compare that data against treaty terms, apply attachment and retention logic, identify recoveries, and flag the exceptions that genuinely need a human.

Act

Agents complete the loop – posting to ledgers, generating statements, routing exceptions – without the hand-off delays that define the manual model.

Process	Today	With AI Agents
Contract ingestion	Manual re-key from slip or cover note	Agent parses, maps to treaty fields, flags discrepancies
Premium allocation	Batch run, manual exception review	Real-time allocation on booking; exceptions auto-routed
Loss recoveries	Claims matched manually	Agent matches, calculates, posts entry, same day
Bordereaux production	Multi-day manual assembly	Agent assembles, validates, distributes on schedule
Close cycle	Days of reconciliation	Agent pre-runs, surfaces exceptions; you approve, not rebuild
Regulatory reporting	Analyst builds Schedule F / Solvency II	Agent populates schedules from live data; analyst reviews

sapiens

What your operation looks like after deployment

From batch to continuous

Right now, ceded reinsurance runs on batch cycles: allocation overnight, reconciliation at month-end, bordereaux assembled period by period. Agents replace that with continuous processing. Policies allocate on booking. Claims match to treaty terms the day they arrive. Month-end close becomes approval-driven – agents pre-run the numbers, surface exceptions, and your team signs off.

Human in the loop, not out of it

More than 90% of routine transactions flow through without a human touch point. Exceptions land on your desk pre-loaded with context: treaty terms, loss history, prior decisions, flag reason. Your senior accountants focus on treaty-edge cases and reinsurer disputes not on building spreadsheets.

Three things your operation needs before deployment

The insurers getting the most from agentic AI led with operating model conditions, not the tool. Three preconditions come up consistently in independent research and you only need them satisfied in the domains where you plan to deploy first, not enterprise-wide.

1. AI-ready data architecture

8 in 10 companies cite data limitations as the primary roadblock to scaling agentic AI. - McKinsey, 2026

Agentic AI amplifies what it finds. For ceded reinsurance, the requirement isn't clean data enterprise-wide – it's data that flows between treaty systems, claims platforms, and the general ledger in the domains where agents will operate first. Organizations that have moved to automated recoverable identification consistently find uncollected amounts in their existing books.

Maps to: Recoverable leakage, reconciliation cost, treaty data latency

2. Codified, executable workflows

Organizations that reimagined workflows before deploying agents consistently outperformed those that focused on the tool. - McKinsey State of AI, 2025

Agentic AI amplifies what it finds. For ceded reinsurance, the requirement isn't clean data enterprise-wide – it's data that flows between treaty systems, claims platforms, and the general ledger in the domains where agents will operate first. Organizations that have moved to automated recoverable identification consistently find uncollected amounts in their existing books.

Maps to: Premium allocation, loss recovery processing, close cycle

3. Governance infrastructure built before deployment

Over 40% of agentic AI projects will be canceled by end of 2027 due to inadequate risk controls. - Gartner, June 2025

Agentic AI acts, not just recommends. The EU AI Act's Article 26 requires human oversight to be structural – built in, not added later. The NAIC's position is equally clear: "Human oversight remains an important part of insurance decision-making," a principle now backed by examination infrastructure in more than 24 states. For ceded reinsurance, this means deciding before go-live who authorizes an agent to post a ledger entry or route a recovery, and how that authorization is recorded in a form regulators and reinsurers can inspect. That framework must exist before agents touch live data.

Maps to: Schedule F and Solvency II compliance, reinsurer audit readiness, AI decision auditability

AI cuts implementation time, phase by phase

Historically, a new reinsurance platform meant 12 months or more from discovery to go-live. Two caveats: AI accelerates field mapping but can't fix bad source data, a quality assessment must come first. And SME commitment through UAT stays high; AI cuts overall client burden by 25-30%, but acceptance testing demands close to 80% availability from senior operations staff.

Phase	Standard	AI saves	What AI does
Discovery and requirements	9 wks	2.5 wks	AI pre-maps requirements from legacy exports
Contract config and migration	14 wks	7.0 wks	NLP parsing of cover notes; 40-60% effort saving on field mapping
Build and configuration	16 wks	3.5 wks	AI-generated user stories and test scripts; parallel workstreams
SIT / UAT testing	10 wks	2.0 wks	Master file validator; AI test packs cut defect triage by ~20%
Go-live prep	5 wks	1.0 wks	Agents co-designed during delivery activate at go-live
Total	**54 wks**	**16 wks saved**	**30% reduction | 38 weeks to go-live | ~20% cost saving**

Where the industry is heading

A Deloitte 2026 State of AI in the Enterprise survey found that 74% of organizations plan to deploy agentic AI within two years, up from 23% the prior year. The organizations using that window to close governance and workflow gaps will arrive at full deployment ready. Those that skip the prerequisite work will move faster on a flawed foundation.

Foundation	Complex treaty management, multi-currency, multi-entity. Bordereaux and regulatory schedules automated.
Phase 1	AI contract ingestion agents. Renewal support tools. Role-based dashboards.
Phase 2	Real-time APIs replace batch feeds. Predictive exhaustion forecasting. ~85% straight-through cash allocation.
Phase 3	Full agentic workflow for ceded operations. Configure without coding. Continuous release cycles.

What enterprise-grade security looks like

Data isolation and PII protection

A PII gateway strips personally identifiable information before data reaches any model. Deployments must operate under enterprise agreements that prohibit logging, storing, or training on your data.

Identity and access control

Agents must inherit, not exceed, the access rights of the user or process they serve. McKinsey's 2026 AI Trust Maturity Survey found only one in three organizations reports mature governance of agentic AI controls – a gap your reinsurers and regulators are tracking.

Auditability and human oversight

Every agent action produces a timestamped audit trail which is a requirement for Schedule F, Solvency II, and reinsurer settlement. Human oversight gates are built in from the start, not added later.

> *In the age of agentic AI, organizations can no longer concern themselves only with AI systems saying the wrong thing – they must also contend with systems doing the wrong thing.*
>
> **– McKinsey AI Trust Maturity Survey, March 2026**

Five questions to start the conversation with your team

Use these to identify where your operation has the strongest foundation for agentic deployment, and where the work needs to start.

1. Reinsurance costs vs. loss ratios

Ask your team: Have costs risen even as loss ratios stayed stable? Treaty pricing may embed a data quality premium your reinsurers are charging but not naming.

2. Workflow documentation

Ask your team: Are treaty interpretation rules, attachment logic, and recovery calculation steps written down and standardized, or do they live with specific people?

3. Recoverable identification

Ask your team: Is it systematic across all open treaties, or does it depend on periodic manual review?

4. Reconciliation time

Ask your team: How many days per period on bordereaux assembly? If it's measured in days, the cost is quantifiable and reducible.

5. AI governance readiness

Ask your team: Have you reviewed your governance framework specifically for autonomous agents, not just generative AI tools?

If two or more answers reveal gaps, you're in the majority. The question is whether you close them before deployment or discover them after. That conversation is where the strongest operations start.

Sapiens brings over 40 years of insurance knowledge across 600 insurers in 30+ countries. With reinsurance, claims and finance connected on one platform, we're ready to move with you and close your gaps.

Talk to our team

sapiens