



Data Readiness for Agentic AI

The Insurance Professional's Guide to
Getting Ready and Staying Ahead

Prospecting



Advising



Renewing




Why Data is What Activates AI

The conversation about [AI in insurance](#) has largely focused on the technology itself, including which models to use, which tools to buy, and which workflows to automate. For most insurance organizations, however, the real unlock is the data itself.

Agentic AI, the kind that autonomously executes multi-step workflows across prospecting, advising, and renewal, is only as capable as the data it can access. Give it complete, connected, well-structured data, and it can act with real autonomy and precision.

The organizations that will get the most out of agentic AI in 2026 and beyond are the ones investing in the data foundation that makes it possible. This e-book is a practical guide to what that foundation looks like, broken down by use case, with the specific data fields and structural requirements that turn AI from a tool you manage into an agent that can act on your behalf.



What Happens Without a Data Foundation

Every insurance organization exploring [agentic AI](#) faces the same fork in the road: invest in the data foundation first, or start building on top of whatever data already exists. The second path is faster to start, and it's also the one most likely to fail.

[Gartner predicts](#) that through 2026, organizations will abandon 60% of AI projects due to a lack of AI-ready insurance data, with 63% either lacking or unsure whether they have the right data management practices in place to support AI at all. An agent asked to flag renewal risk or benchmark a client's coverage can only work with what is visible, and when employee counts are outdated or coverage data lives in three disconnected systems, the output looks complete while resting on partial information.

↓ 60%

abandonment of AI projects due to lack of AI-ready insurance data

↓ 63%

lack adequate data management practices to support AI

That first disappointing experience carries its own cost. Producers who had a rough run with an early AI tool are slower to trust the next one, and leadership starts asking harder questions before approving future investments. Meanwhile, agentic AI capability keeps compounding, and organizations that delay are widening the distance between where they are and where the market will be, making each future catch-up attempt more expensive than the last.

None of this is a reason to slow down on AI. Investing in a data foundation is the fastest, highest-leverage step an insurance organization can take right now, and it's the one that determines whether every future AI investment compounds or collapses. As [McKinsey recently highlighted](#), insurance sector AI leaders have created 6.1 times the total shareholder return (TSR) of AI laggards by investing in tools that enable a strong data foundation.

What Makes Data "Agentic-Ready"?

That foundation isn't one-size-fits-all, though, and not all data problems are equal. Some gaps slow AI down, while others make it unreliable in ways that aren't obvious until a client is on the other end of the output. Here's the framework for assessing where your data stands today.

Complete — Every account, contact, and policy record has the fields that matter populated. Partial records create partial outputs. Agentic AI doesn't guess; it works with what's there.

Current — Stale data produces stale results. Renewal dates, contact information, coverage details, and exposure data need to reflect the actual state of the account, not what was entered at binding two years ago.

Consistent — Data entered differently across producers, offices, or systems breaks the logic agentic workflows rely on. Consistent classification, naming conventions, and field usage are prerequisites for reliable AI output.

Connected — Data that lives in silos, such as across an AMS, a separate CRM, spreadsheets, and carrier portals, can't be accessed by an AI in a single workflow. Integration and connectivity determine what the AI can reach.



Master Data by Use Case

Prospecting (Producer Agent)

Agentic prospecting requires the AI to identify, research, and reach out to high-fit prospects without producer involvement at every step.

The master data required:

- **Prospect firmographic data:** Company name, industry classification (NAICS/SIC), employee count, annual revenue, location
- **Risk signals:** Industry-specific exposure indicators, loss history proxies, operational characteristics
- **Contact data:** Decision-maker name, title, verified email, LinkedIn profile
- **Outreach history:** Prior touchpoints, response status, do-not-contact flags
- **Producer assignment:** Which producer owns the territory or relationship

Data gap most commonly seen:

Contact data quality. Outdated or missing decision-maker emails result in outreach going to the wrong person or bouncing entirely.

Renewal Management (Advisor Agent)

Agentic renewal management requires the AI to scan the full book, score accounts by risk and opportunity, and surface a prioritized action queue 90 days out.

The master data required:

- **Policy data:** Expiration dates, carrier, premium, coverage lines, endorsements
- **Account data:** Client name, industry, size, years with agency, relationship tier
- **Loss history:** Claims filed, open claims, loss ratios by account
- **Engagement data:** Last contact date, last meeting date, open service requests
- **Coverage gap data:** Known gaps or underinsured exposures flagged at prior renewal

Data gap most commonly seen:

Loss history and engagement data. Many agencies track policy data well but have inconsistent records of client interactions and claims history.

Master Data by Use Case

Client Advising and Benchmarking

Agentic advising requires the AI to generate relevant, industry-specific benchmarking summaries and talking points for client meetings.

The master data required:

- **Industry classification:** Accurate NAICS/ SIC codes at the account level
- **Coverage structure:** Current limits, deductibles, and coverage lines in structured fields (not PDFs)
- **Exposure data:** Payroll, revenue, fleet size, locations — updated annually
- **Benchmarking reference data:** Industry peer data for comparison (typically provided by platform intelligence layer)

Data gap most commonly seen:

Coverage structure stored as documents rather than structured fields. AI cannot read a PDF the way it reads a database record.

Carrier Submission

Agentic submission workflows require the AI to compile, format, and route submission packages without manual assembly.

The master data required:

- **Applicant data:** Legal entity name, FEIN, contact information, years in business
- **Exposure data:** Complete and current for the relevant line of business
- **Loss runs:** Structured, current, and attached to the correct account record
- **Prior carrier and premium history:** At least three to five years
- **Supplemental application data:** Completed and current for the target market

Data gap most commonly seen:

Loss runs not attached or outdated. This is the single most common reason agentic submission workflows require human intervention.

Perspectives From the Experts: Tech Team Interview



**Data quality is the foundation
of every successful AI initiative.**



Sherrie Clevenger

*Vice President of Product Management,
Zywave*

Sherrie Clevenger, Vice President of Product Management at Zywave, noted that the most common data issues typically manifest in the following ways: “Duplicate and outdated client or policy records top the list, often from years of manual entry across multiple systems. Missing or incomplete fields are next, especially expiration dates, coverage details, and carrier information that never got filled in because no workflow required it at the time.”

Clevenger also spoke to the underlying reason behind these data gaps, saying “Agencies typically have data spread across the AMS, a separate CRM, and spreadsheets that never sync, so there is no single source of truth. Notes and attachments are frequently unstructured (free text in comment fields, PDFs that were never indexed) which an agent cannot reliably read or act on. Finally, naming and classification conventions vary by producer or office, so the same line of business or account type gets entered differently depending on who touched it.”

Resolving these problems, and ultimately, enabling growth across the board, comes down to consistency, data hygiene, and process improvement guidance. “The best-structured customers treat their AMS as the single system of record and keep a consistent taxonomy across producers and offices, rather than letting each person define fields their own way,” Clevenger said.

“Zywave’s dedicated data experts help customers assess, cleanse, and optimize their data to meet the requirements of AI and agentic workflows. Through a combination of data analysis, gap identification, and process improvement guidance, we help organizations enhance data readiness, accelerate adoption, and maximize the impact of AI-driven automation and decision-making.”


The AMS as the Agentic Foundation

Agentic AI operates inside a connected stack and at the center of that stack for insurance organizations is the agency management system.

The AMS is where the master data lives: accounts, policies, contacts, loss history, renewal dates, coverage details. When that data is clean, complete, and consistently structured, the AMS becomes the fuel that powers every agentic workflow. When it isn't, every downstream AI output reflects those gaps.

Zywave's AMS is built with this in mind. Rather than functioning as a passive system of record, it serves as an active data layer, one designed to integrate with the Zywave Apex MCP server and surface the right data at the right moment in every agentic workflow. Producers don't have to manually pull records or copy data between systems. The AI reaches into the AMS, finds what it needs, and acts.

That connection, between a well-structured AMS and an MCP-enabled agentic platform, is what separates agencies that use AI from agencies that are transformed by it.



Why Your Data Vendor Matters as Much as Your AI Vendor

As insurance organizations build out their data foundation, most of the attention goes to choosing the right AI tools. But the vendor behind your data is just as important; arguably, it's the more foundational decision.

A data partner built specifically for insurance brings clean, connected, industry-specific data that agentic AI can act on from day one. With deep, purpose-built data assets across the insurance distribution lifecycle, Zywave gives organizations the kind of connected foundation that turns agentic AI into a working advantage.

Your Data Readiness Checklist	
Account and Contact Data	
All active accounts have complete legal entity names and contact records	☐
Primary decision-maker contact information is verified and current	☐
Industry classifications (NAICS/SIC) are populated for all commercial accounts	☐
Accounts are consistently named across all systems	☐
Policy and Coverage Data	☐
Renewal dates are populated and accurate for all active policies	☐
Coverage lines, limits, and deductibles are stored in structured fields	☐
Endorsements and coverage changes are reflected in current records	☐
Prior carrier and premium history is available for key accounts	☐
Engagement and Activity Data	☐
Last contact date is tracked at the account level	☐
Open service requests are logged and current	☐
Outreach history is recorded and associated with the correct account	☐
Loss and Claims Data	☐
Loss runs are attached to account records and current	☐
Claims history is structured and queryable (not only in PDFs)	☐
Loss ratios are available for renewal scoring	☐

Where to Start

Data readiness doesn't have to be an all-or-nothing project. The most effective approach is to prioritize by use case and to align your growth plan with the appropriate solution. Zywave's Sherrie Clevenger noted that "**Growth intent** should shape the recommendation, not just the agency's current pain point."

An agency focused on growing its book benefits most from pairing Advisor Agent or Producer Agent capabilities with tools that support new business and cross-sell, such as prospect nurturing and rating and account rounding workflows, because the value shows up as new premium and expanded accounts.

Clevenger said that, "An agency focused on **efficiency or retention** is better served by leaning into service automation and renewal workflows, where the value shows up as time saved and reduced errors rather than new business generated."

The agencies that get the fastest results from AI-ready insurance data are the ones who know where their gaps are and address them systematically, use case by use case, workflow by workflow.



Ninety-five percent of the time should be spent out and about, not sitting behind a computer. If companies don't have a tool like this, they are still going to spend a lot of time gathering contact information, pulling leads, and then structuring and creating emails.



Chris Whitney

*Partner,
Sterling Benefits*

Sterling Benefits is an early adopter of Producer Agent, part of the Zywave Apex™ AI platform. Producer Agent pulls together prospect and contact data, drafts outreach sequences, and runs them automatically, so producers are not spending the bulk of their time on manual work.

Ready to activate agentic AI in your agency?

The first step is understanding where your data stands. Talk to a Zywave expert about your data readiness and how Zywave Apex, and the AMS built to power it, can help you get there by building a strong foundation of AI-ready insurance data.

[Learn More About Zywave Apex](#)

