



**The Service Moment
Insurers Cannot Automate
Incorrectly: *A Diagnostic
for Claims and Contact
Center Leaders***



Table of Contents



1	How the AI conversation in insurance service has changed from the “what” to “how”	03
2	Why the insurance service center is the most challenging area	04
3	Hurricane surges: The problem carriers don’t put on their roadmap	05
4	Six questions to consider	06–07
5	Distinct operational risks depending on your specific carrier model	08
6	The flywheel most carriers are not closing out	09
7	Defining operational capabilities	10
8	Bring this diagnostic into your next planning room.	11



How the AI conversation in insurance service has changed from the “what” to “how”

A year ago, the primary question in most insurance carriers' conversations would be “*what should we be doing with AI?*” That question is long answered.

What has replaced it is a harder question. Insurance carriers now have something working – be it in the form of an AI pilot, a proof of concept (POC), or a use case that survived the committee and made it to production. Behind all these, there is a queue of ideas that nobody can get through.

The latest 2026 statistics present a grim picture. While **90% of carriers** have run AI in at least one department, only 7–22% have fully scaled their AI systems to production.

This is not an AI model problem, but an operating model problem, and it shows up first (and probably the worst) in the service center, because this is where call volume, regulatory exposure, and customer sentiment land in the same queue.

Meanwhile, the solution that most carriers are still receiving is about cloud migration, revolving around:

- Move to the cloud.
- Modernize the data warehouse.
- Consolidate the data estate.

For a lot of carriers, that work is already done or is in progress. They have built a digital file cabinet. What they cannot do yet is safely act on what is in it. While most cloud migration is about consolidating data, the competitive advantage is now in the orchestration of those interactions.

This eBook presents six crucial questions about **AI in insurance** service operations, along with the “honest” answers that tell you whether you are ready to scale anything.

Why the insurance service center is the most challenging area

Insurance is a trust-based business currently navigating a service transformation. At its core, the carrier's primary mandate is twofold: the secure stewardship of the policyholder's contract, and the delivery of promise when that contract is invoked. The service center is the point where those two mandates converge—or collide.

- Keeps a very large file cabinet.
- Runs a call center that services a contract.



That call center performs the following completely different jobs:



Routine servicing

This is high volume, patterned, and transactional. Here's an example of routine servicing: "I bought a new car. My daughter turned sixteen. I am adding a motorcycle to the insurance policy." Effectively, it is the natural starting point for automation. Most carriers know this and have already (or started to) automate this.



First Notice of Loss (FNOL)

This is the initial report that customers submit in the event of loss, damage, or theft of their insured assets. Here's an example of FNOL: "There is a tornado, and I need somewhere to sleep tonight. My car is in a ditch. There was a fire." Same queue, same staff, same phone number, but completely different stakes.

FNOL is the single most important service moment for any insurance company. This is when the customer finds out the "real" value of their insurance policy. The brand reputation of the insurance provider can either be "confirmed or destroyed" in that customer call.

For most insurance carriers, claims processing (not customer service) is a critical moment. Here is why that matters for **insurance claims automation**. When routine servicing and FNOL run through the same operational path, any automation applied evenly across both jobs will eventually apply a routine pattern to a crisis call. That is the failure in **insurance automation** that can cost a carrier a customer, a complaint, and sometimes even the regulator's attention. It is also entirely predictable and preventable, but in reality, most carriers don't separate the paths before starting.



Hurricane surges: The problem carriers don't put on their roadmap

No carrier can staff for a hurricane surge.

An additional headcount to absorb a catastrophe can escalate their expenses even in normal periods. Instead, when a hurricane storm crosses three states, their call volume goes up, and the overflow is assigned to an external party or a contracted center.

In this heightened state, your brand's most critical service interactions are being delivered by third-party staff who must have been onboarded for your specific book of business within days, rather than weeks. Most AI pilots are architected for the 'normal Tuesday'—they fail to account for this operating reality, where the weakest link is an improvised, externalized workflow that has yet to be stress-tested against your unique compliance and handling standards.

Every carrier knows this, but few AI plans account for it.

If your **contact center AI solution** does not have an answer for hurricane surges, it's designed only for normal days.





Six questions to consider

What matters is not the actual answers, but how quickly and confidently they get answered, because that is the real measure of whether your operation can absorb any future surge.

Here are six diagnostic benchmarks (in question format) that can help you assess your operational readiness:



Question #1.

Who answers the customer's insurance claim call during a catastrophe?

If the answer is a business partner, follow up on how long it takes them to bring them up to speed. What decisions are they allowed to make, what do they do with their learnings, and are you responsible for those decisions after the contract expires?

Question #2.

What percentage of first customer calls are resolved without a transfer?

This is not for the current quarter but for daily. If the number takes a week to produce, you do not have an AI observability problem in your current operation. In this scenario, any level of **insurance automation** will be invisible in exactly the same way.





Question #3.

If you deploy an AI assistant in the service center tomorrow, who is allowed to change what it says?

This is the question that ends most conversations. For a carrier with thousands of employees, the honest answer is usually nobody. Not IT, because it's a customer-facing language. Not the business, because it is a system. Not compliance, because they were not asked.

Question #4.

Who signs off on that change, and how long does it take?

Following up on the above question, if changing a line of script takes weeks, your automated customer engagement platform & service center will be permanently out of date with actual policies, insurance products, and regulatory posture.

Question #5.

During claim processing, does any learning from the customer interaction reach the product pricing team?

Most carriers say "yes," meaning the data exists somewhere. That's not the same as underwriting in a form and on a timeline that changes next year's hypothesis.

Question #6.

Can you show a regulator how a decision in a workflow was made?

This is not something you could explain, but show it, with a record, for a specific case, on a specific date. If the answer requires a person to reconstruct it from their memory, that is the ceiling on how far automation can go in a regulated function.

The strategic implications of your answers vary by market position. Here is how your response profile changes based on your risk exposure and business model.





Distinct operational risks depending on your specific carrier model

The six questions are universal because FNOL is structurally the same everywhere. What the answers mean is not.

If you are a mutual service provider, your policyholders are your owners, which makes the service moment the closest thing to brand value. It also means commitment to get FNOL right with the least surplus budget to spend. Question #4 is the one to watch. If a change to customer-facing language has to travel the same governance path as a capital decision, you will never keep an automated service layer updated. The fix is not more speed in the boardroom. It is a defined, lower-threshold path for operational language changes, decided once, in advance.

If you write long-tail lines, workers' compensation and medical professional liability do not have claims; they have long-lasting relationships where claims adjusting longevity serves as a core differentiator. Once a file opens, it stays open through the treatment, litigation, reserve changes, and settlement. Everything written about AI in claims assumes a faster auto claim that opens and closes in a week, and almost none of it applies to you. Your answer to question #5 is not whether the claims data reaches underwriting, but whether the file's own history is legible to the adjuster who inherits it after three years and two handoffs. Continuity, not cycle time, is your metric.

If you are cat-exposed. For a Florida-based property writer or a regional carrier in a wind or wildfire state, surge is not an edge case but the operating model. You already know that your worst week is your least-controlled week and that the third party answering your phones during a storm is also learning your book in real time. Question #1 is your whole exam. For anything that lets catastrophe capacity be absorbed with your rules, your language, and your record is worth more than any efficiency gain in an ordinary quarter.

If you distribute rather than underwrite, wholesale brokers, retail agency networks, and motor clubs have the inverse problem of a carrier. You own the relationship, but don't own the claim. Your service moment is advocating for a customer inside an external process or timeline, with limited visibility. Question #6 is more important for you than for a carrier, as the concerned record is spread across parties. Question #5 has a twist: you see the customer but rarely see the outcome, which means the most valuable feedback loop in your business is the one you are least able to close.

If you are benefits- or health-adjacent, high-frequency, lower-severity claims at scale can change the math. Volume makes automation obviously attractive and makes small error rates expensive, fast. Your risk is not a mishandled crisis call but a systematic error repeated thousands of times before being noticed. Question #2 can protect you. If you cannot see resolution and exception rates in real time, it's difficult to catch a pattern failure early enough to be affordable.



The flywheel most carriers are not closing out

Step back from the customer service center for a second and consider the overall business loop.

Underwriting sets a hypothesis about risk management. Insurance policies get written against it, followed by claims. What actually happened gets compared to what was assumed to happen, thereby improving the following year's premium pricing structure.



That loop constitutes the carrier's primary economic engine, yet for most, it remains decoupled. Claims data lands in systems that underwriting struggles to parse, in formats that are inconsistent, and on timelines that rarely align with the pricing cycle. The result is a pricing hypothesis refreshed by aggregate industry data and 'gut feel' rather than the carrier's immediate, claim-level experience. Fixing this isn't just about data integration; it's about bridging the cultural and temporal gap between the claims floor and the underwriting desk.

When managed as a cost exercise, service center modernization produces a cost outcome - Shorter handle times, lower cost per claim, and an improving headcount line. When managed as a data exercise, modernization turns the flywheel. When clean, structured, and captured-at-source, claims data is just what underwriting needs. The service center is not just where you spend money but where the company is losing its most valuable data.



Defining operational capabilities

A resilient, modern insurance service center demonstrates six essential operational capabilities:



Routine and high-stakes paths are separated.

Adding a vehicle and reporting a total loss do not travel the same road. This is an operating decision before it is a technology one, and it is the most affordable on the list.



Surge is absorbed, not outsourced blindly.

Catastrophe capacity exists without the third party learning your business from scratch every time and without your worst week being your least-controlled week.



Someone owns what the system says.

A named person, not a committee, with authority to change customer-facing language.



Change is measured in days.

Any customer updates from the service layer take less time than updating the customer policy.



Every interaction is reconstructable.

Not summarized, but reconstructable, case by case, for a regulator, whenever demanded.



The claims data reaches underwriting on time.

The flywheel closes.

Each of these 6 characteristics is operational at the core. However, every carrier is missing at least three of them.



Bring this diagnostic into your next planning room

We have packaged these six questions into a clean, one-page workshop tool designed to help your leadership team align on where your service-center vulnerabilities lie.

Download the 1-Page Diagnostic Worksheet <https://docs.google.com/presentation/>

Request a facilitated 15-minute briefing with our Insurance team to help walk through the results.
<https://www.onixnet.com/financial-services/>



 onixnet.com
 connect@onixnet.com
 800.664.9638

Get in touch

Follow us:



Copyright © 2026 Onix . All rights reserved.

