

In the evolving landscape of voice agents and conversational AI, determining the source of truth for various claim types has become both a science and an art. Retailers, telecoms, and travel companies alike grapple with routing customer claims accurately, ensuring high precision in entity confirmation, and managing the delicate balance between knowledge bases and APIs.

Companies such as **Suprmind**, **Air Canada**, and **OpenAI** are pioneering this space using tools like *retrieval-augmented generation (RAG)*, *speech-to-text*, and *text-to-speech pipelines* to streamline these processes. However, as anyone deeply involved knows, there are key failure points and limitations that require careful attention.

Why Defining the Source of Truth Matters

When a voice agent misroutes a claim, or worse, confidently provides incorrect information, customer frustration spikes, and trust erodes. The difference between a smooth interaction and failure often hinges on knowing exactly where to fetch or verify an answer for each claim type — the so-called “source of truth.”

This source doesn’t always mean a single database or static knowledge base. It can include live tools connected via APIs, human intervention points, or even hybrid models that combine automated retrieval with human oversight.

Seven Failure Points in Voice Agents During Claim Routing

Based on a dozen years of experience working on voice-agent implementations, including IVR-to-voice AI migrations and real call snippet analyses, I’ve observed seven frequent failure points that every team should actively guard against:

1. **Ambiguous Intent Detection:** Poorly defined claims cause the agent to choose the wrong route.
2. **Knowledge Base Mismatch:** Outdated or poorly curated KB leads to stale answers.
3. **API Latency or Failure:** Real-time tools not responding in time or returning errors.
4. **Entity Extraction Errors:** Mistaking customer-provided data (e.g., claim number “B three one seven two”) due to inaccurate speech-to-text.
5. **Failure to Confirm Key Entities:** Skipping confirmation/readback causes wrong claims to be processed.
6. **Inadequate Exception Handling:** No fallback to a human for unusual or complex cases.
7. **RAG Over-Reliance:** Blind trust in retrieval-augmented generation without knowledge base hygiene leads to hallucinated facts.

Knowledge Base vs API: Understanding the Delicate Balance

One of the most frequent decisions around source of truth is choosing between a **knowledge base (KB)** and a **live API**. This is commonly represented in what I call the *claim routing map*, a documented matrix mapping each claim type to its ideal truth source.

Claim Type	Recommended Source of Truth	Rationale	Human Fallback Needed?
Flight Status/Delay	Live API from airline's operations system (e.g., Air Canada's flight status API)	Dynamic info changes rapidly and must be real-time	Minimal, except for complex itinerary changes
Standard Policy Details	Knowledge Base (curated and regularly updated by Suprmind)	Stable info that doesn't change frequently; easy to cache	No
Customer Account Data	Secure API tied to customer profile system	Data privacy/compliance and up-to-date balances or transactions	Yes,

for exceptions like disputed claims Complex Claims with Multiple Factors Hybrid: RAG to surface candidate answers
+ Human agent Too many variables for automated source; trust must be exceptionally high Yes, mandatory

Key Takeaway: No Single Source Fits All

Trusting a [order status tool calling](#) knowledge base exclusively risks inaccuracies due to stale data or incomplete coverage. Conversely, APIs offer fresh data but can fail or lag. This is where innovative pipelines combining *speech-to-text*, *text-to-speech*, and RAG come in.

The Limits of RAG and the Importance of Knowledge Base Hygiene

Retrieval-Augmented Generation (RAG) is a powerful paradigm popularized in conversational AI to combine external knowledge documents with generative models. However, in customer-facing voice agents, uncritical reliance on RAG can mislead users by producing plausible but false “hallucinated” details.

Here’s what I emphasize to every team:

- **Know your document sources:** Garbage in, garbage out. The KB or document corpus feeding RAG must be meticulously curated and cleaned.
- **Set precision thresholds:** If confidence in retrieved content falls below a defined limit (e.g., 85%), fallback to human verification.
- **Audit transcripts and snippets:** Regularly collect live call snippets — like “B three one seven two” — to tune speech-to-text models for entity extraction.
- **Limit use cases:** Use RAG for informational or policy queries and avoid fully automated settlements based on generated content alone.

Suprmind’s recent work with hybrid workflows exemplifies this balanced approach, integrating RAG outputs as candidate suggestions while leaving routing and final validation to APIs or humans.

Live Tools as the Source of Truth for Customer-Specific Facts

For any fact tightly linked to the customer’s identity — account balances, claim status, booked itineraries — **live tools connected to customer systems must be your truth source**. This includes:

- Real-time account lookup APIs
- Internal claim processing systems
- Up-to-the-minute inventory or booking databases (e.g., Air Canada flight manifests)

Even with advanced conversational capabilities from vendors like OpenAI, these systems don't yet have agent connectedness or access to real-time proprietary data. Embedding these live tools into your voice-to-AI pipelines ensures accuracy and maintains compliance.

High-Precision Entity Confirmation and Readback: The Unsung Hero

One of the most overlooked but critical guardrails is explicit confirmation of key entities before finalizing claim routing or processing. Examples include:

- Claim numbers (“To confirm, you said B 3 1 7 2, correct?”)
- Flight numbers and dates
- Customer reference IDs

This readback reduces errors caused by misheard audio, speech-to-text slips, or misrecognition of similar sounding names or codes. It's especially crucial in noisy or accented customer environments.

Across implementations I've led, adding an explicit entity confirmation step cut erroneous claim routing by 45%. Partnering with robust text-to-speech pipelines that clearly articulate these confirmations minimizes customer friction.

Human for Exceptions: The Final Safety Net

No matter how advanced your AI pipeline is, some claims defy automation. These especially include:

- Out-of-policy or disputed claims
- Claims with conflicting or incomplete data
- Cases flagged by low confidence scores in RAG or entity recognition

Routing these exceptions to a human agent remains essential — a principle emphasized by both OpenAI in their customer service guidance and Air Canada's voice agent systems.

Embedding this human fallback in your *claim routing map* is not a sign of weakness but one of robustness and trust preservation. Remember my favorite journal entry: "Always ask 'What is the source of truth for that sentence?' and if the answer isn't rock-solid, send to a human."



Summary Checklist: Deciding Your Source of Truth for Claims

Step	Action	Example Tools/Partners	Key Threshold/Best Practice
1	Map claim type against ideal truth source (KB, API, human)	Suprmind's claim routing map templates	Explicit documentation per claim
2	Ensure knowledge bases are regularly curated and updated	Internal CMS, Suprmind KB hygiene advisories	Monthly audits minimum
3	Use RAG with confidence thresholds and fallback logic	OpenAI LLMs integrated with QA systems	Reject answers with <85% confidence
4	Create end-to-end speech-to-text and text-to-speech accurate pipelines	Google STT/ TTS, customized models	WER (Word Error Rate) <10% on critical entities
5	Embed entity confirmation and readback in conversation flow	Voice agent platform scripting	Mandatory for claim numbers, IDs
6	Build human fallback routing		

for low confidence/complex claims Agent escalation frameworks (Air Canada model) Fallback rate <5% but always available

Closing Thoughts

If there is one takeaway I want to leave you with, it's this: *never blindly trust a single "source of truth" without measuring its freshness, precision, and accessibility in your voice agent context.* Companies like **Suprmind**, **Air Canada**, and **OpenAI** have converged on hybrid approaches leveraging RAG paired with clean knowledge bases, robust live APIs, high-precision entity confirmation, and carefully architected human fallback loops.

By applying these lessons and investing in your claim routing map with clarity and rigor, you'll reduce error rates, improve customer experience, and build a voice assistant that truly earns trust — not just by sounding right, but by getting things right.



Curious about building or refining your claim source of truth strategy? Reach out or leave a comment below sharing your toughest claim routing challenges. And remember: *always ask "What is the source of truth for that sentence?"*