taxnav™

×

PQ Impact

Why Most AI Chatbots Fail in Regulated Industries

And what we did instead — a behind-the-scenes look at automating support for TaxNav

A PQ Impact Case Study

The Problem: When Growth Becomes a Support Crisis

In April 2026, the UK government begins enforcing Making Tax Digital for Income Tax (MTD IT) — requiring self-employed individuals and landlords earning over £50,000 to submit their tax records digitally to HMRC on a quarterly basis. By 2028, the threshold drops to £20,000, bringing millions more into scope.

TaxNav is an HMRC-recognised MTD IT software platform. It handles quarterly submissions, year-end final declarations, spreadsheet bridging, and CIS deductions — all through a browser-based interface aimed at making compliance as frictionless as possible.

As MTD IT deadlines approach, TaxNav's user base is growing rapidly. And with that growth comes an inevitable challenge: support volume is scaling faster than the team. Tax software isn't like a typical SaaS product. Users don't ask "how do I change my password." They ask things like "do I need to include rental income from a jointly owned property in my Q2 submission?" The stakes are real — a wrong answer could mean an incorrect tax filing.

How do you scale support for a product where every answer has to be correct, and the subject matter is complex and regulated?

Why Standard AI Chatbots Don't Work Here

The obvious approach is to connect an LLM — OpenAI, Anthropic, or any other provider — via API, feed it your documentation, and let it handle user queries. For most SaaS products, that's reasonable. For tax software, it's a liability.

1 The hallucination problem is not a minor inconvenience

LLMs generate plausible-sounding text — but they don't know whether what they're saying is true. In tax compliance, a chatbot that confidently tells a user they don't need to declare certain income is a legal risk. We needed every answer traceable to a verified source. Zero room for fabrication.

2 Tax data is sensitive data

TaxNav users enter income figures, property details, CIS deduction records. Many AI solutions involve sending queries to external infrastructure, creating a data sovereignty problem. We needed an architecture where user data never leaves a controlled environment.

3 Support queries are domain-specific and nuanced

"How do I add a second property?" sounds simple. But the correct answer depends on ownership structure, income thresholds, and which quarter the user is filing for. An LLM — even one using RAG — would get this dangerously wrong.

The Comparison

These constraints eliminated every off-the-shelf approach we evaluated.

Requirement	LLM-based approach	AI4U approach
Answer accuracy	Probabilistic — may guess	**Deterministic — verified knowledge base**
Data privacy	Data sent to external APIs	**Data is private, securely encrypted, and not accessible to anyone else.**
Domain specificity	General or fine-tuned	**TaxNav-specific Q&A pairs**
When AI can't answer	Fails silently or guesses	**No hallucinations, says when it doesn’t know**

Which meant we had to build something different.

What We Built Instead

Three design decisions shaped our approach — each a direct response to the constraints above.

01

A verified knowledge base, not a language model generating answers

AI4U doesn't "think." It matches user queries against a curated database of question-answer pairs built for TaxNav. Every answer is verified before going live.

If no matching answer exists, the system doesn't guess — it escalates. In a regulated domain, "I don't know" beats confident fabrication.

The knowledge base is auto-generated from TaxNav's existing documentation. No manual authoring, no prompt engineering required.

02

Self-contained architecture — no data leaves the system

AI4U processes all queries within its own infrastructure. User data, conversation context, query content — none routed to external APIs.

For tax compliance handling sensitive financial data, this is a hard requirement — not just a preference.

Designed for data sovereignty from day one, not retrofitted as an afterthought.

Intelligent Escalation —
The AI Knows When to Step Back

03

AI4U monitors the entire conversation and detects when a user needs human help — regardless of whether the AI technically "knows" the answer. Frustration, complexity requiring judgement, need for human reassurance — these are cases where information alone isn't enough.

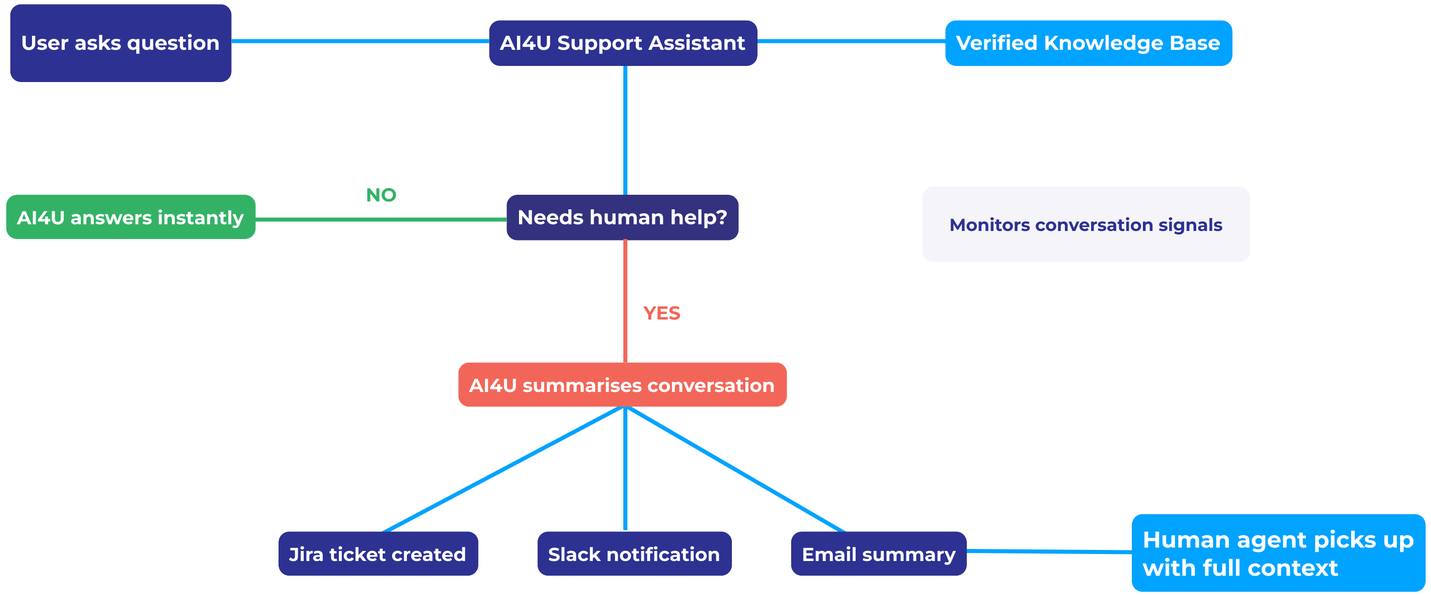

The human agent doesn't start from scratch — they get the conversation summary, user's problem, and full interaction history.

What This Means Beyond TaxNav

If you're building any product where the cost of a wrong answer is high, the same constraints apply.

1. "Just add AI" is not a strategy.

Define your failure mode first. If a wrong answer means legal risk, financial loss, or eroded trust — you need a different architecture than a general-purpose LLM.

2. Privacy can't be an afterthought.

If your product handles sensitive data, your AI solution needs data sovereignty from day one. Retrofitting privacy onto external-API architectures is expensive and fragile.

3. The best AI support makes humans faster, not redundant.

AI4U resolves the majority of queries alone. But its real value shows in escalation: automated handoff, conversation summary, Jira ticket before the agent opens the case.

4. Your existing docs are more valuable than you think.

AI4U's knowledge base is auto-generated from existing help articles. No manual authoring, no prompt engineering. If you have decent docs, you're closer than you expect.

Thinking about AI for your product?

Whether it's support automation, onboarding, internal tools, or something else — the challenges are the same: accuracy, privacy, and making sure AI helps rather than creates new problems.

We designed and built the entire AI4U system for TaxNav — from R&D and architecture through to implementation and project management.

It's what we do: we take complex technical problems and turn them into working software.

Get in touch →

PQ Impact — Software Agency & AI Solutions
pqimpact.com | ai4u.digital | marcin@pqimpact.com