Chapter 1: High-Stakes Interactions: Why They're Rewriting the Rules of Customer Support

Customer Support Complexity Is Intensifying

Despite 50+ years of evolution, customer support still struggles with fragmented systems, rising expectations, and high-emotion scenarios. McKinsey reports that 92% of consumers were dissatisfied with service in 2023—highlighting how traditional automation hasn't kept up with real-world demands.

What's driving the dissatisfaction? Well, just like most answers – it's a combination of factors.

1. **Customers expect resolution, in real time, on any channel:** Customers expect seamless service—on any channel, at any time. Consumer tech has raised the bar, and now all industries are expected to match that level of digital care.
2. **We sell more products & offer more services:** Support complexity has grown. A simple refund request might trigger fraud checks, localization logic, regulatory validation, and backend handoffs.
3. **Complicated issues translate into emotionally charged interactions**: Consumers have been trained to expect a resolution that is seamless, simple, and fast. Those expectations are more aggravating when dealing with emotionally urgent matters such as personal finances, insurance accessibility, or even a broken consumer electronic.
4. **These strenuous experiences are not limited to customers:** Customer support agents are strained by staff turnover, talent shortages, limited training, and juggling multiple systems. It's no wonder that a support agent finds it difficult to address the more common but complex customer challenges.

These emotionally urgent issues—like financial disputes or insurance access—raise expectations even higher. And agents, juggling multiple systems and rising workloads, are rarely equipped to meet them.

While AI handles up to 80% of routine inquiries, it still fails on tasks that demand cross-system context, regulatory precision, or emotional nuance. Only 19% of customers trust bots with sensitive issues—and 77% of agents now spend more time on complex interactions.

So, what makes something complex? And why are these situations such high stakes for enterprise companies today?

High-Stake vs Routine Tasks: Understanding the Difference

High-stakes tasks aren't just more complicated—they carry consequences. Mistakes can trigger financial penalties, reputational damage, compliance breaches, or even safety risks.

These scenarios increasingly becoming more and more common in customer service, demanding a higher degree of control over how they are handled:

Criteria	Routine Tasks	High-Stakes Tasks
Workflow complexity	Few steps, linear	Multi-step, cross-system workflows
Emotional urgency	Low	High—emotionally or financially sensitive
Knowledge required	Basic, standardized	Nuanced, often judgment-based
Error impact	Low risk, reversible	High risk, often irreversible
Systems involved	Few, often self-contained	Multiple systems with real-time dependencies
Team coordination needed	Rare	Frequent—cross-functional collaboration needed
Examples	Password reset, FAQs	Claims, refunds, warranty disputes

The most common element distinguishing the high-stakes complex request from the more mundane is the degree of contextual information & human judgment required to resolve the issue. These are issues where customers are asked a variety of questions to assess how best to move forward. Even more than personally identifiable information, context can change and therefore the agent must make judgement calls as to how to move forward.

Take Susan, who works for a nationally recognized Credit Union. Her team manages acquisitions, or more commonly known as loan applications. As part of the loan application process, customers must endure a complex series of questions about their financial situation.

A single keystroke error during a loan application derailed the entire process. Despite a 30-minute intake, the issue escalated to IT—too late to retain the customer, who had already moved to a competitor. In workflows like this, the tolerance for error is near zero.

Multiple calls and exchanges later, the customer had already left to try their luck with a competitor Credit Union. Imagine where it happens with more than one application – leading to a financial loss in loan approval rates. The tolerance for error here was next to zero. Further, acknowledging the customer's frustrations with empathy could have further reduced the risk of exit.

High-stakes cases demand a very low error rate and a high degree of emotional intelligence. Routine cases, by contrast, often just need to be handled quickly and correctly – they are perfect candidates for automation or AI **so long as accuracy is assured**.

Customers are Turning Away From Self-Service

Despite technological advancements, companies have not been able to successfully automate high-stake tasks through self-service. The industry-standard approach of using AI chatbots and self-service search tools simply doesn't cut it.

Customers are frustrated with self-service AI. Many chatbots have become nothing more than a workaround to reach a live agent faster—especially when live chat isn't available.

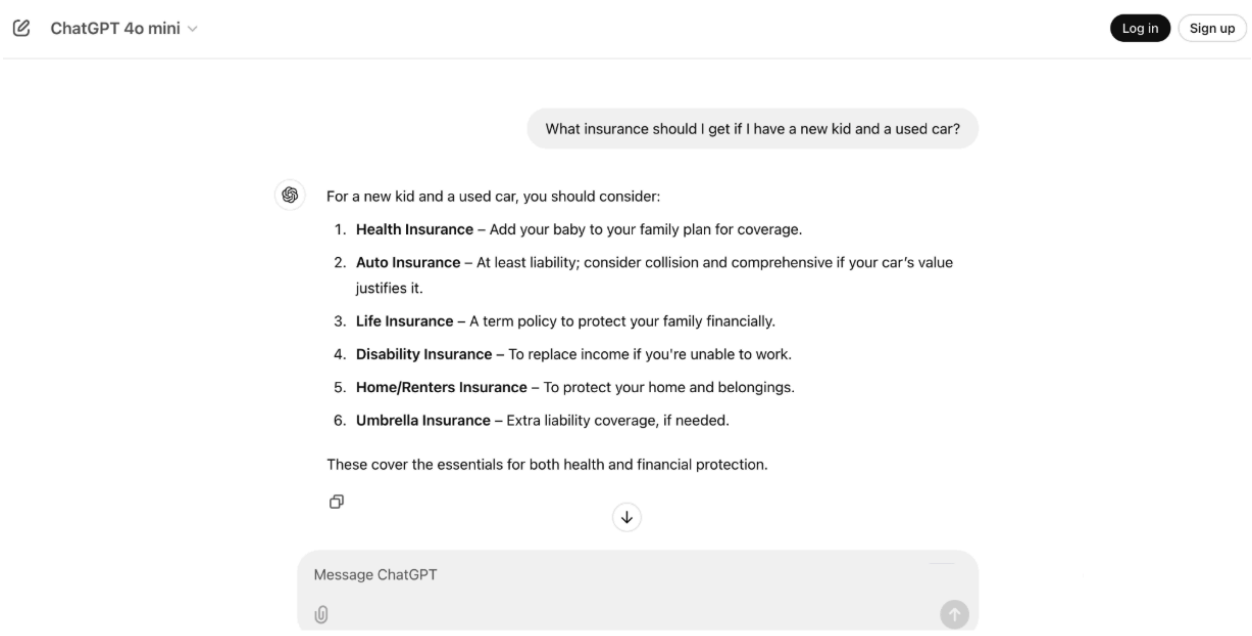

When it comes to complex issues, customers bypass bots—because AI lacks access to real-time customer history, status, product availability, and eligibility rules.
When customers encounter complexity, they turn to human agents. But even then, the issues remain.

Complexity Persists—Even With Human Agents

When complex queries reach human agents, the challenges don't disappear. Most support agents have multiple tabs open. IVR, CRM, inventory, and homegrown systems-just to let customers know if their claim is approved or denied.

This screen-hopping kills resolution times and results in a lack of control over their own support processes. Instead of solving problems efficiently, agents are forced to:

- Navigate fragmented systems
- Piece together information from disconnected sources
- Ask customers to repeat details they've already provided

The lack of unified context leads to repeated questions, inconsistent responses, and "start-over" moments. This leads directly to prolonged interactions, lower First-Call Resolution—increasing operational costs by up to 15% per contact—and greater likelihood of errors.

Despite agent involvement, complex scenarios often remain unresolved due to inadequate tools and a lack of unified context.

The expectation is clear: 70% of customers assume support agents have full access to their complete context. Yet, only 22% of organizations report having unified customer data.

The Real Cost: Revenue Loss, Churn, and Poor CSAT

The fallout from poorly managed complexity is severe. It directly contributes to declining revenue and diminished CSAT scores.

30% of customers abandon brands after just one negative chatbot interaction.

73% of customers cancel ongoing purchases if support fails mid-interaction.

Ineffective handling of complex customer support issues also leads to agent burnout and higher attrition rates. The impact isn't just external. Replacing one agent costs ~$10,000. For teams over 250, annual attrition can exceed $1M in lost productivity and training.

To solve complex support challenges, companies must go beyond AI for routine tasks.

Has your customer support organization identified the high-stakes vs. routine customer service issues on your remit? Take a moment to consider your own processes and:

1. **Explore your customer journeys: Which** activities could cause real harm to your customers? Where could human error be more easily handled by machine? Which repetitive or routine tasks are your agents over-indexed on?
2. **Review your guardrails & policies:** Does regional localization influence certain standard protocol? Are there legal ramifications associated with entering incorrect information? Identify the regulated decisions that could benefit from increased oversight and assess whether it's better handled by AI or a human?
3. **AI Adoption & Brand Trust:** What's the current state of customer sentiment in your brand and in your industry as a whole? What is the level of AI adoption by your target customer audience? Many industries, such as healthcare, are reticent to invest significantly in AI because their own customers are less trusting of AI accuracy. Is that the same for your business?

Chapter 2 will explore why the industry as a whole has not yet conquered complex workflows and how major barriers—ranging from AI maturity to rigid, short-sighted chatbot design—keep solutions trapped at the routine level.

White Paper: Chapter 2

I thought we were already using AI? Why hasn't AI made this any easier?

Think of today's support stack like a high-tech car navigating crumbling infrastructure. The tools are smarter—but the environment is chaotic, with outdated systems and unpredictable inputs.

Customer service has followed a similar path.

Despite a decade of NLP advances, most automation still relies on rigid scripting. A few decision trees and help articles can't resolve emotionally charged or multi-system issues.

But just like our modern cars can't prevent traffic jams, our upgraded support tools haven't solved for **complexity**. In many cases, this is due to a lack of control over how the AI is applied to those complex situations

Yes, our automation systems have gotten more powerful. But at the same time, our customer issues have become more layered — more emotional, more situational, more dependent on cross-system data and judgment. And that's where traditional automation starts to fall short.

Why Agentic AI?

The short answer – customer expectations have outpaced our ability to deliver, even with the AI available.

Early chatbots did their job — they surfaced answers to common questions and deflected basic inquiries from call centers. But they were never designed to untangle a billing dispute, resolve a broken device claim, or guide someone through a high-emotion insurance appeal. These are the types of situations where the stakes are high, emotions are elevated, and resolution paths aren't always clear.

Over time, support leaders began to realize something important: **The biggest friction isn't in the tools themselves — it's in the fragmented, ever-changing context that surrounds the customer issue.**

That's where Agentic AI comes in.

What is Agentic AI?

Agentic AI represents the next stage in advancement of machine learning. We went from simple algorithms, to predictive algorithms, and now with Agentic AI, we have trained models to perform specific tasks and in some cases, make decisions on our behalf. They're intelligent agents that can:

- Pull data from multiple systems
- Interpret and reason through steps in real-time
- Make decisions based on business rules and user context
- And, when needed, escalate to a human — with all the relevant background in hand.

The rapid deployment of AI Agents has become a possibility in the last couple of years thanks to a few converging trends:

1. **GenAI Maturity -** Large language models like Gemini, GPT4, etc. have a far better understanding of natural human communication.
2. **Cloud-native architecture-** modern systems can connect, extract, and share information with disparate data systems like CRMs, inventory systems, and knowledge basis all at once.
3. **Economies of scale:** this one isn't new. Support teams are asked to do more with less.

Forrester analysts note that *agentic AI is a "breakthrough capability" poised to become a competitive necessity,* as it enables **self-directed decision-making and execution** beyond what traditional automation or analytics could do.

How will Agentic AI meet the needs of complex workflows?

Agentic AI is designed to handle the layered, situational nature of complex support workflows—where resolution depends on both systems and human-like judgment.

Recall, some of the key characteristics of complex workflows:

1. Activities requiring multiple handoffs, between systems & people. These feature the extraction of information from multiple systems and depending on human judgement to interpret said information.
2. Activities associated with emotional urgency - both on the part of the customer as well as the agent working to address the issue
3. A dependency on context - or as sometimes referred to as situational awareness.

Agentic AI systems have some of these technical requirements in place:

1. They can automate decision making, even with workflows requiring multiple sequences
2. They can extract critical operational & product data in real-time
3. They can interpret data and make decisions based on these interpretations

Traditional bots are like road trips that require printed maps. Agentic AI is GPS—it reads conditions in real time and keeps the journey smooth.

Who Should Build the Agent?

Internal knowledge experts are your best AI architects. They understand nuance, maintain institutional logic, and ensure decisions reflect both compliance and customer context.

Agentic AI does represent the next phase of customer service automation technology - taking the basic desire for clear and action oriented decision workflows – applying AI, and transforming them into Agentic systems that can streamline and simplify these workflows.

Involving your internal knowledge experts as AI Agent authors ensures more control over the output - and ultimately, more control over the AI technology itself.

Chapter 3: Building AI Agents That Work: Zingtree's Control-Driven Approach

I. The Challenge: Why Control is Essential for Effective AI Agents

Susan's story illustrates a central truth: even small data entry errors can derail high-stakes workflows. You need AI that's accurate—and under your control.

The truth is, most AI solutions weren't built for Susan's world.

With Agentic AI on the rise, leaders must prioritize tools that provide clarity, control, and alignment with real support environments.

Before jumping into the Agentic fray - a few considerations that Support Leaders must consider. The same as evaluating any technology:

1. Garbage in, Garbage out: We've always known that missing details, lack of access to updated information, or just plain inaccuracy can result in poor resolution of customer issues.

2. Design of Agentic AI offerings: one other basic tenet of technology investments – look to subject matter expertise. Does the vendor you are evaluating have the credibility to address your decision trees? Do they understand your business? Do they understand how many different systems you need to manage as an agent?

Ultimately, it's not about having an AI agent -it's about having one where you retain control over the guardrails and having an AI agent that you can trust.

Meet Zingtree - the AI Agents you can Control

Zingtree is the only platform that enables support teams to build their own AI Agents - by giving them control over the logic, rules, and context that shape every customer interaction.

Built by teams with deep support experience, Zingtree evolved from rule-based workflows to a platform that combines low-code design with GenAI-driven autonomy.

We built Zingtree to accommodate how teams engage with AI. Businesses that operate in significant regulatory environments or service customers dealing with private and secure information, adapting to AI technology can take time. Digital transformation and process transformation are not always lean or agile. Zingtree was built to accommodate the stop/go nature of digital transformation. The platform harnesses both low-code and GenAI features to enable teams to generate strict, rule based processes with limited AI integration – all the way to the autonomous agentic features that can easily recognize customer intent & apply these to the right workflow, at the right time.

Zingtree integrates with your key systems, transforms data into action, and puts human oversight at the center of agent development.

As one VP of Customer Support said, "Our agents are living in a manual tax form. We need to move them to TurboTax." That shift only works if AI understands your business—and you define the logic behind it.

Zingtree flips the AI script. Instead of relying on black-box models to guess intent, it lets your team define the logic, inputs, and guardrails. AI handles the delivery. You stay in control.

What that looks like in practice:

- **Controlled inputs**: You choose what data the agent sees—and what it doesn't.
  Built-in compliance: Every step is traceable and aligned with your internal standards.
- **Real-time context**: Zingtree connects to CRMs, policy engines, and inventory systems to deliver accurate, moment-specific guidance.
- **No-code agility**: Operations teams can build and update flows without IT.

This input-first approach puts your internal experts in the driver's seat. You don't just deploy AI—you direct it.

Why Control Matters Now More Than Ever

As AI adoption accelerates, the real differentiator is control. Zingtree makes automation operational—not aspirational.

Zingtree is the answer for teams that want to automate with confidence. That means:

- Scaling without adding headcount
- Shortening onboarding and resolution times

- Reducing errors and repetitive escalations
- Improving agent satisfaction and retention
- And future-proofing your tech stack without ripping it apart

As Bill put it: *"Ultimately, where I really want to go is that only exceptions hit our support team. Everything else should be handled through self-service—accurately, automatically, and in line with how our business works."*

With Zingtree, that vision isn't just aspirational. It's operational. Zingtree can help you automate everything accurately so that only exceptions are handled by your support team and everything else is automated through self-service.

Conclusion: Build the Agent You Can Trust

Agentic AI is powerful. But in complex support environments, power without precision is a liability. Zingtree gives your team the tools to build AI agents that deliver results—not just responses.

You don't have to trade control for speed. You can have both.

Let your agents focus on what matters. Let your customers get answers faster. And let Zingtree help you build automation that works—for everyone.