LiveHelpNow

Customer Service Operations *Playbook*

The seven steps I take every support team through before they buy software, add AI or hire another agent.

Michael Kansky
Founder, LiveHelpNow. In business since 2010.

BEFORE YOU START

Fix the operation first. *Then* buy the tools.

Most support teams shop for software to fix a problem they have not measured yet. This playbook is the order I work in: find out what customers actually ask, where each answer lives, what a solved problem really costs, and only then decide what a bot, an agent or a new hire should take on.

It is written for the person who owns support: a head of support or CX, an operations lead, or the owner of a company where support still lands on their desk. Each step takes days, not months, and each ends with a decision you can act on.

01 Get access to the knowledge

02 The diagnostic before any vendor evaluation

03 Size the real cost of support

04 Decide what the bot handles and when a person takes over

05 Staff to the hour, not the average

06 Run the weekly review

07 Three mistakes this playbook exists to prevent

Figures from outside sources are named where they appear. Planning assumptions are marked as starting points: replace each one with your own numbers as soon as you have them.

STEP 1

Get access to the knowledge

The hardest part of fixing support is rarely the software. It is getting permission to see what your own company already knows.

Before anyone sorts a single conversation, collect access to every place the answers and the evidence live. In most organizations that means getting through some red tape: legal, IT, privacy and the department that "owns" each system all have to say yes.

Customer conversations

Chat transcripts, email and ticket history, social messages, survey comments.

Phone transcripts

Call recordings or transcripts, voicemail, callback notes.

Knowledge bases and wikis

Help center articles, internal wikis such as Confluence, shared drives, PDFs.

CRM and account data

Account records, orders, billing, the fields an agent looks up on every call.

Admin SOPs

Refund rules, escalation paths, approval limits, the "how we really do it" documents.

The people

The two or three agents everyone asks. Their heads hold what no system does.

> **Do this:** make one list of systems, the owner of each, and who has to approve access. Ask for read-only access and a data handling agreement up front. A single approval meeting with legal and IT beats ten separate requests.

Without this step, every later number is a guess, and every vendor demo runs on the vendor's data instead of yours.

STEP 2

The diagnostic before any vendor evaluation

Pull your last 90 days of conversations and sort them by reason.

Start with your top 20 reasons and measure what share of your volume they represent. For each reason, write down three things: how many contacts it drives, how long it takes to handle, and where the answer lives.

The answer lives in one of four places:

Where the answer lives	What it means
A written document	AI can answer it well, if the document is current.
The customer's own account record	AI can answer it, once the customer's identity is verified.
A person's head	Write it down first. No vendor can automate what nobody has written.
A judgment call	Stays with people.

That last column decides everything. Answers in documents and account records are what AI handles well. Answers in someone's head have to be written down before any tool can help. Judgment calls stay with your team.

> **Some reasons should not exist at all.** If 15 percent of your chats are "where is my order," the fix may be a better shipping email, not a bot. Fix those at the source before you automate anything.

Output of this step: a one-page table of your top 20 reasons, with volume, handle time and the answer's location for each. It is the brief you hand every vendor.

STEP 3

Size the real cost of support

Cost per ticket flatters you. Cost per contact is better. Cost per resolution is the honest number: a problem solved that stays solved.

Cost per resolution = total fully loaded monthly support cost ÷ issues resolved and not reopened within 7 days

What goes in

- Wages, payroll taxes and benefits
- Software: help desk, chat, phone and AI
- Management, quality assurance and training time
- Recruiting and ramp time for replacements
- Outsourcing or after-hours fees
- Credits and refunds issued to fix service failures

What most teams forget

- **Shrinkage.** An agent paid for 40 hours is not on the queue for 40 hours. Meetings, training, breaks and time off take a real share. Allow roughly 25 to 35 percent when planning, then replace the assumption with your actual figures.
- **Attrition.** Every replacement costs weeks of training and reduced productivity before they work on their own. Calculate the actual ramp time from your last few hires; for ordinary support, 8 to 12 weeks is a fair budgeting assumption, longer for technical or regulated work.
- **Repeat contacts.** One problem that takes three contacts is counted three times as work and once as a result.
- **The internal tax.** The time engineering, finance and the owner spend answering support's questions. It never shows up in the support budget.
- **Customers who gave up.** Not a cost line, but it is lost revenue.

> **When you compare to AI:** the major vendors charge roughly $0.99 to $2 per resolution. Ask exactly what they count as resolved. Some count a customer who simply left.

STEP 4

Decide what the bot handles and when a person takes over

The bot handles questions where the answer is in an approved document or in the verified customer's own record, and nothing needs deciding.

An agent handles exceptions to policy, money beyond a set limit, complaints, upset customers, anything legal, medical or safety related, and cancellations of valuable accounts.

A handoff must happen when

Trigger	Rule
The customer asks for a person	The first time. No "are you sure?"
The bot misses twice on the same question	Two is the safe limit. Genesys found 84 percent of consumers give a virtual agent three tries or fewer.
The bot cannot find a source	It says so and hands off. It never guesses.
Identity cannot be verified	On any account question.
The topic is on the never-list	Legal, medical, safety, complaints, cancellations of valuable accounts.
A refund or credit is outside policy	Any refund or credit outside a documented, pre-approved policy goes to a person. Until that policy exists, the bot issues none.

> **The handoff always carries the full conversation,** so the customer never repeats themselves. A route to a person is not a failure of the bot. Gartner found 87 percent of customers say a route to a human is essential.

Sources: Genesys, State of Customer Experience, 2026 (5,811 consumers in 20+ countries). Gartner, 2026.

STEP 5

Staff to the hour, not the average

Coverage hours

Chart contacts by hour of the week for the last 6 to 8 weeks. Staff to the hour, not to the daily average. The average hides Monday at 10am.

Headcount

Take volume per hour times handle time, run it through an Erlang C calculator to hit your response-time target, then add shrinkage. Plan for two simultaneous live chats per agent; three may work for straightforward questions, so measure quality before you raise the limit. Voice is one at a time.

After hours

Decide by what an unanswered contact costs you overnight. If it is a sale or an urgent account problem, staff it, in house or with an overnight team. If it can wait until morning, use AI plus a clear promise, such as "a person will reply by 9am." Staff nights when the cost of making customers wait consistently exceeds the cost of coverage.

> **A quick break-even:** an 8-hour night shift at $30 an hour costs $240. If a contact answered overnight is worth $40 to you, the shift pays for itself at 6 contacts a night. This is a financial check, not a staffing forecast: response targets and minimum coverage still apply.

Callback or live queue

Offer a callback when the expected voice wait passes about two minutes, or when abandonment starts climbing, provided you can honor it. Callbacks keep the customer's place in line. A live queue wins when the problem is blocking them right now: they cannot log in, a payment is failing, or they are in the middle of a purchase.

> **A callback you do not honor on time is worse than a queue.** Only offer what your schedule can keep.

STEP 6

Run the weekly review

Thirty minutes, once a week, eight numbers. Each one drives a specific decision.

Number	Decision it drives
Top 10 contact reasons, change vs. last week	Which to fix at the source, which to automate next
AI resolution rate by reason (not reopened in 7 days)	Expand or pull back AI on that reason
Handoff rate, and why	Content to write, rules to adjust
Questions the AI could not answer	This week's writing assignments, each with an owner
Repeat contact rate	Whether "resolved" is real
First response and resolution time vs. target, by hour	Schedule changes
Cost per resolution, human vs. AI	Budget and headcount
Satisfaction on AI vs. human conversations	Where customers are being let down

> **The rule:** any number that does not drive a decision gets dropped from the review.

STEP 7

Three mistakes this playbook exists to prevent

I have watched each of these happen. Each one cost the company more than the software did.

> The clients below are anonymized. Their names can be disclosed, and they can serve as references, on request.

1. Automating a problem that should not exist

At Client 1, one question kept coming back to support week after week, and the company was ready to buy a chatbot to absorb it. When we read the conversations, the question existed because of a gap in the company's own process: customers were not told what they needed to know at the moment they needed it. We fixed the message at the source first, and the volume dropped before any bot was involved.

What it cost: agent hours every week spent answering a question that should never have been asked, and a software purchase that would have automated the gap instead of closing it.
What I learned: before you automate a question, ask why customers have to ask it.

2. Counting a conversation as resolved when the customer was still stuck

At Client 2, the dashboard looked good: most conversations were closing without an agent. When we read a sample of those conversations, many customers had not been helped at all. They left the chat, then came back by email or phone with the same problem.

What it cost: repeat contacts the dashboard counted as wins, and agent time spent cleaning up after the bot.
What I learned: a conversation ending is not a problem solved. Measure whether the customer has to contact you again.

3. Staffing to the daily average instead of the hours customers need help

At Client 3, the daily numbers said the team had enough agents. The hourly chart told a different story: a few peak hours of the week carried much of the volume, while staff were spread evenly across the day. Customers waited longest exactly when most of them showed up.

What it cost: long waits and abandoned contacts at peak, overtime to catch up, and quiet hours with little to do.
What I learned: customers do not arrive at the daily average. Staff for the hours they need you.

ONE PAGE

The checklist

- [ ] Access approved to conversations, phone transcripts, knowledge bases, CRM and SOPs, read-only, with one data handling agreement.
- [ ] Last 90 days of conversations sorted by reason; top 20 listed with volume, handle time and where the answer lives.
- [ ] Reasons that should not exist identified, each with a fix at the source.
- [ ] Answers that live in people's heads assigned to someone to write down.
- [ ] Cost per resolution calculated with shrinkage, attrition, repeats and the internal tax.
- [ ] Bot scope written: approved documents and verified account records only.
- [ ] Handoff triggers written, including the refund policy and the never-list.
- [ ] Contacts charted by hour of the week; staffing set to the hour.
- [ ] After-hours and callback rules decided and kept.
- [ ] Weekly review on the calendar with eight numbers and an owner for each.

Run this playbook on LiveHelpNow

LiveHelpNow puts chat, tickets, phone and the knowledge base in one workspace, with Hue, the AI that answers from your own content. Hue learns from your knowledge base, past chats and tickets, your website and the documents you upload; it shows its source, says when it is not sure, and hands off to a person with the full conversation. A knowledge gaps check lists the questions it cannot answer yet, which is Step 6's writing list.

Your data is used to answer your questions. It is not used to train AI unless you choose to allow it.

Book a demo

livehelpnow.net · 877-548-3001 · 436 N Main St #1107, Doylestown, PA 18901