FREE MEMBER RESOURCE

AGENCY GUIDE

The 24/7 AI customer service setup

Launch always-on AI support inside HighLevel — for your clients and your own agency. Faster answers, fewer late-night messages and a service you can package, price and resell.

Charles Higgins · Nexus Hub

Where HighLevel agencies get unstuck.

CONTENTS

What's inside

HOW TO USE THIS GUIDE

Every chapter works on two levels: set this up for a client account and charge for it, or run it inside your own agency's sub-account. Wherever the two diverge, we call it out.

00 — INTRODUCTION

The opportunity hiding in the inbox

Every local business you serve has the same problem: leads and customers expect immediate answers, at any hour, on whatever channel they happen to reach for. The owner can't be available around the clock — and neither can you, for every client at once.

That gap is where agencies get stuck fielding "did you see my message?" texts — and it's also one of the cleanest recurring-revenue offers you can build. HighLevel already ships everything you need: Conversation AI to answer across chat, SMS and social; Voice AI to catch the calls nobody picks up; workflows to route, follow up and escalate. You don't need to evaluate a market of chatbot tools. You need a repeatable setup.

This guide walks you through that setup end to end: auditing where a business actually loses money to slow replies, configuring the AI to sound like the brand, launching safely, keeping humans in the loop and turning the whole thing into a packaged service you can deploy from a snapshot in an afternoon.

Why 24/7 service matters — and why it's your sale to make

- Speed wins the deal. The first business to respond usually gets the job — a lead who messages three plumbers at 9pm hires the one that answers at 9:01pm.
- Every unanswered message is churn fuel. Slow replies erode the trust your marketing built.
- Your clients' competitors are automating. Always-on response is becoming the baseline, not the differentiator.
- For you, it's sticky revenue. A client whose phones, texts and DMs are answered by a system you run does not casually cancel.

> FROM THE FIELD
>
> One agency deployed Conversation AI for a home-services client who woke up to dozens of unanswered messages every morning. Within a month, the AI was handling overnight inquiries and booking estimates straight onto the calendar. Sales rose 28% — not from more leads, but from finally answering the ones already coming in. The agency now sells that exact setup to every new client.

Let's turn customer service from the thing that burns out your clients — and your team — into the most defensible line on your invoice.

01 — DISCOVERY

Find the pain points before you build

Don't open the bot builder yet. The setups that fail are the ones configured on assumptions. The ones that stick are built on a week of real conversation data — and that data doubles as your sales pitch.

Step 1: run a conversation audit

For one week, review everything landing in the client's Conversations tab — SMS, email, Facebook and Instagram DMs, Google Business Profile chat, webchat — plus their call log. If the client isn't in HighLevel yet, this audit is your onboarding week. For each inquiry, track:

- The question asked
- Time received and time answered
- Channel it came in on
- Simple or complex
- Whether it needed contact-specific data (order status, appointment time, account details)

EXERCISE — INQUIRY TRACKING TEMPLATE

RECEIVED	CHANNEL	QUESTION	TYPE	RESPONSE	AI-READY?
Mon 8:15am	GBP chat	"What are your hours?"	Simple	3 hours	Yes
Mon 11:30pm	Instagram	"Has my order shipped?"	Simple + data	9 hours	Yes
Tue 2:45pm	Phone	"It's not working, I've tried..."	Complex	10 min	No

After a week, pull four numbers: the share of questions that are simple and repetitive, the average response time by channel, the volume arriving outside business hours and the top five question types. Most local businesses find 60–80% of inquiries are routine — the AI's job — and only 20–40% need human judgment.

Step 2: put a dollar figure on the status quo

This calculation closes the deal, so do it properly. Estimate three costs:

- **Time cost** — hours per week the owner or staff spend answering questions, times their hourly value.
- **Opportunity cost** — leads lost to slower competitors. Count the after-hours inquiries that never got a same-day reply.
- **Stress cost** — the owner answering DMs at dinner. Unquantifiable, and often what actually signs the contract.

EXAMPLE MATH

A client spending 15 hours a week on inquiries, at $50/hour of owner time, is burning $750 a week — $39,000 a year — on answering questions. If AI handles 70% of them, that's 10.5 hours and $525 a week reclaimed. Suddenly your $400/month AI-support retainer is the cheapest employee they've ever hired.

Step 3: rank the highest-impact questions

From the audit, shortlist the questions that come up most often, take the longest to answer, arrive after hours or sit closest to a sale. These are what you train first. For most local businesses the list looks like:

- Hours and location
- Availability and booking
- Refund and cancelation policies
- Payment methods
- Pricing and quotes
- Order or job status
- Service-area coverage
- Basic troubleshooting

For your own agency: run the same audit on your agency inbox. "How do I reset my password," "can you resend the invoice" and "is the campaign live yet" are AI-ready questions too — and answering them by hand is why your team never gets to billable work.

02 — THE STACK

Your stack is already in HighLevel

The original version of this playbook — the one written for business owners — spends a chapter comparing chatbot vendors. You get to skip it. Everything below ships inside the platform you already pay for, works in every sub-account and can be packaged into a snapshot.

The five pieces, and what each one does

CAPABILITY	WHAT IT HANDLES
Conversation AI	Two-way AI replies across SMS, webchat, Facebook Messenger, Instagram DMs, Google Business Profile chat and live chat — trained on the business's own content
Voice AI	Answers inbound calls, fields common questions, captures caller details and books appointments when nobody picks up
Workflows	The routing layer: missed-call text back, after-hours autoresponders, escalation to humans, follow-up sequences, review requests
Knowledge base training	Feed the bot the client's website, FAQs and policy docs so answers come from their content, not the model's imagination
Snapshots	Your productization tool: capture the whole build once, deploy it to the next client in minutes

Step 1: decide what to turn on first

Resist launching everything at once. Sequence by where the audit showed the most bleeding:

- **Missed calls are the top complaint** — start with missed-call text back plus Conversation AI on SMS. It's a one-day build with instant, visible wins.
- **Web leads go cold overnight** — start with the chat widget plus Conversation AI, wired to the calendar.
- **The phone rings off the hook** — start with Voice AI as the after-hours and overflow answerer.

Step 2: check the integration points

The AI is only as useful as what it can see and do. Before you build, confirm in the sub-account:

INTEGRATION CHECKLIST

- Calendars are connected and availability is current, so the AI can actually book
- Social channels (Facebook, Instagram, GBP) are linked in the sub-account's integrations
- Phone numbers are provisioned and A2P registration is complete — no registration, no SMS delivery
- Pipelines and custom fields exist for what the AI should capture (service type, urgency, address)
- Contact data the AI needs to reference — order status, membership level — lives in fields it can read

Step 3: build once, snapshot everything

Do your first build in a template sub-account, not directly in a client account. Use custom values for everything client-specific — business name, hours, service area, booking link, escalation phone — so the bot prompts and workflows read from variables instead of hard-coded text. When the build works, save it as a snapshot. Client two takes an afternoon: load snapshot, fill in custom values, retrain the knowledge base, test, launch.

That snapshot is the difference between selling your hours and selling a product. Guard it accordingly.

Train the AI to sound like the brand

An AI that answers correctly but sounds like a robot reading terms and conditions will get your client complaints instead of compliments. Voice is a configuration problem, and it starts with a worksheet — one you'll fill out in the kickoff call for every client.

Step 1: document the brand voice

EXERCISE — BRAND VOICE WORKSHEET

- If this brand were a person, how would customers describe them? (friendly, professional, playful, no-nonsense)
- Which three to five adjectives define the communication style?
- Are there greetings or phrases the business uses consistently?
- What does it call its customers — clients, guests, members, patients?
- Contractions? Exclamation points? How formal is too formal?
- What topics must the AI never improvise on? (medical claims, legal advice, guarantees, allergens)

The answers become the personality section of your Conversation AI prompt. A boutique bakery might land on: warm, friendly and slightly playful; calls customers "friends"; uses contractions; never casual about food safety or allergens. Paste exactly that into the bot's instructions — the more specific the worksheet, the less generic the bot.

Step 2: write the answers to the top questions

Take the high-impact list from chapter one and draft a model answer for each, in the brand's voice. Keep answers complete but short, link out for detail and always end with a forward motion — a question or a booking offer.

MODEL ANSWER — "WHAT ARE YOUR HOURS?"

"We're open Monday–Friday 9am–6pm and Saturdays 10am–4pm. Closed Sundays to recharge! If you need something outside those hours, I can take a message or book you in for our next opening — anything specific you're planning to come in for?"

Answers the question, sounds human, covers the after-hours case and keeps the conversation moving toward a booking.

Step 3: map response trees for multi-step scenarios

Some inquiries branch. Map the branches before you configure the bot's goals, so it collects the right things in the right order. The classic example is a status inquiry:

- **Opening move:** "Happy to check on that — what's the order number, or the email you used?"
- **Has the number:** confirm the order, give status and expected date, offer tracking.
- **Has only an email:** disambiguate between recent orders before answering.
- **Has neither:** don't guess. "To protect your privacy I'll hand this to our team — want someone to chat now or email you?" Then escalate.

Step 4: build the knowledge base

Point the bot's training at the client's website and upload what the site doesn't say. The test: could a smart new hire answer customers using only what you've fed the bot? If not, keep feeding.

EXERCISE — KNOWLEDGE BASE PRIORITY LIST

List the top 20 things the AI must know cold, each as a short plain-language explanation:

- Each core service and what it includes
- Refund and cancelation policies
- Service areas or shipping zones
- Current promotions
- Pricing structure and payment options
- Warranty and guarantee terms
- Booking process, step by step
- What to do in an emergency or urgent case

Snapshot note: keep the voice worksheet and knowledge list as onboarding forms in your snapshot. Client intake fills them; you paste, train and test. Repeatability is the product.

04 — LAUNCH

Set up the launch

Configuration is where good bots go bad. A careful launch sequence protects the client's brand — and your reputation as the agency that installed the thing.

Step 1: run a closed beta before going live

- Start the bot in suggestive mode: it drafts replies, a human approves each one before it sends
- Recruit five to 10 testers — your team plus the client's staff — and have them attack it with real questions, badly phrased
- Log every wrong, awkward or off-brand answer and patch the knowledge base or prompt
- Only flip to auto-pilot on a channel once the bot clears a full test week without an embarrassing miss

Testing always surfaces phrasings you didn't plan for. Cheaper to find them in beta than in a client's Google reviews.

Step 2: configure when and how the AI engages

- **Trigger timing** — does webchat proactively greet after 30 seconds, or wait for the visitor?
- **Operating window** — 24/7 with different messaging after hours, or after-hours only as a human backstop?
- **Introduction** — the AI names itself and is honest about being an assistant. Transparency builds trust; pretending fails eventually and publicly.
- **Handoff criteria** — when does the conversation go to a human? (Next step covers this.)
- **Follow-up** — a workflow nudge if a visitor abandons mid-conversation: "Still around if you need anything!"

EXAMPLE INTRO MESSAGE

"Hi! I'm Maya, the virtual assistant for [Business Name]. I can help with questions about services, scheduling and pricing — and our team reviews every conversation. What can I help you with?"

Step 3: build clear paths to a human

No AI handles everything, and the fastest way to enrage a customer is trapping them in a bot loop. Wire escalation triggers into your workflows from day one:

- Customer asks for a human, in any words — hand off immediately, no questions
- Two to three failed answer attempts — escalate automatically
- Complaint language or frustration — route straight to a person
- Sensitive topics — payments, disputes, personal data, anything legal or medical
- Requests needing judgment or an exception to policy

On escalation, the workflow should assign the conversation to a named user, fire an internal notification and add the contact to a follow-up pipeline so nothing waits until someone happens to check the inbox. Set a response-time expectation in the handoff message and hold the client's team to it.

Step 4: handle privacy and compliance

- Complete A2P 10DLC registration before any SMS automation goes live — unregistered traffic gets filtered
- Honor opt-outs everywhere; keep required stop language in first-touch SMS
- Know where conversation data lives and set retention expectations with the client
- Check the rules that apply to the client's industry and region (GDPR, CCPA, HIPAA-adjacent caution for anything health)
- Publish a plain-language disclosure on the client's site

DISCLOSURE TEMPLATE

"Our AI assistant helps us respond faster. Conversations may be stored and reviewed to improve our service, and are never sold or shared for marketing. You can reach a human anytime — just ask."

Balance AI and human support

The systems that keep clients for years aren't the most automated ones — they're the ones where AI and humans each do what they're best at, and the handoff between them is invisible.

Step 1: draw the line, in writing

AI OWNS

- Factual answers — hours, prices, coverage
- FAQs and policy questions
- Collecting contact and job details
- Booking and rescheduling appointments
- Status checks
- Simple product or service recommendations

HUMANS OWN

- Complaints and upset customers
- Complex troubleshooting
- Policy exceptions and refunds beyond policy
- Personalized advice
- Price negotiation
- VIP relationships

Put this table in the client's SOP. When their staff knows the bot's lane, they stop fighting it — and stop apologizing for it.

Step 2: make handoffs seamless

The cardinal sin of escalation is making the customer repeat themselves. HighLevel keeps the full thread in Conversations, so there's no excuse:

- The human reads the thread before typing a word
- They open by proving it: "Hi Sarah — I've read your conversation about the installation issue, and I see you've already tried resetting it. Let's try something else."
- Escalations carry a priority tag — a refund dispute outranks a rescheduling question
- Pause the bot on that thread once a human steps in, so the two never talk over each other

Step 3: install a review rhythm

The bot gets better only if someone reads what it got wrong. Build this into your retainer — it's deliverable work, not overhead:

IMPROVEMENT CHECKLIST

Weekly

- Review 10–15 escalated conversations; find the patterns
- Add three to five new answers to the knowledge base
- Update anything stale — prices, hours, promotions

Monthly

- Pull resolution rate, escalation rate and satisfaction trends
- Rework the conversation flows with the highest escalation rates
- Re-test the top 20 questions; retrain where answers drifted
- Steal phrasing from the best human replies and fold it into the bot

Step 4: train the client's team to work with it

A one-hour training session at launch saves months of friction. Role-play the core moves: taking over a conversation from the bot, spotting a knowledge gap and reporting it to you, and knowing when to let the bot finish versus stepping in. Set the expectation clearly — escalations get a human response within the promised window, every time.

For your own agency: the same rhythm applies to your client-support bot. Your account managers review its escalations weekly, and it learns your clients' recurring questions the same way.

Measure it, report it, price it

The metrics chapter matters twice as much for an agency: the numbers that prove the system works are the same numbers that justify the invoice and win the next client.

Step 1: capture the baseline before launch

You already have most of this from the chapter-one audit. Lock it in writing before the bot goes live: average response time, first-contact resolution rate, weekly hours spent on inquiries, after-hours inquiries answered (usually zero) and inquiry-to-sale conversion. No baseline, no before-and-after story.

Step 2: track the AI’s own numbers

- **Resolution rate** — conversations closed without a human
- **Escalation rate** — and the reasons, grouped
- **Bookings made by the bot** — the metric clients quote to their friends
- **After-hours volume handled** — work that previously just didn't happen
- **Accuracy** — spot-check answers against the knowledge base monthly

METRIC	BEFORE AI	WEEK 1	WEEK 4	WEEK 12
Avg response time	5.2 hours	2 min	1.8 min	1.5 min
First-contact resolution	65%	72%	78%	85%
Customer satisfaction	4.1/5	4.0/5	4.3/5	4.6/5
Weekly staff hours on inquiries	15 hrs	8 hrs	6 hrs	4.5 hrs
After-hours inquiries handled	0%	92%	95%	97%

Step 3: calculate ROI — theirs and yours

THE CLIENT'S ROI

Hours saved × hourly value × 52. Ten hours a week at $50/hour is $26,000 a year — against a $400/month retainer ($4,800/year), a 5.4× return before counting a single after-hours booking. Add recovered revenue: bookings the bot made outside business hours are sales that used to walk.

YOUR ROI

Rebill the AI usage with margin, or wrap it into a flat tier. Once the snapshot exists, each new deployment is a few hours of configuration against recurring retainer revenue — and the monthly review rhythm from chapter five is billable, recurring and genuinely valuable. This is the productized-service math that makes agencies scalable.

Don't skip the qualitative wins in your reporting: the owner sleeping through the night, consistent answers on every channel, capacity that scales with demand. Those lines renew contracts.

Step 4: send a monthly report that sells the renewal

One page: conversations handled, resolution rate, bookings made, after-hours saves, hours returned to the team and what you improved this month. Every report is a small re-close. Use the data to optimize too — knowledge gaps become new training content, stuck flows get simplified and confusing answers get rewritten.

07 — EXPANSION

Expand what the AI can do

Once the core setup runs clean for a month, expansion is where retainers grow. Each capability below is an upsell conversation backed by the data you're already reporting.

Step 1: go from reactive to proactive

Support that reaches out first feels like magic and drives revenue. Build these as workflows feeding the same AI conversation layer:

- **Abandoned booking recovery** — started the calendar flow but didn't finish? Text them 20 minutes later.
- **Post-purchase check-ins** — confirm delivery, share usage tips, catch problems before they become reviews.
- **Reorder and maintenance reminders** — consumables running low, seasonal service due.
- **Review requests** — timed to the moment satisfaction peaks, with the AI fielding any "actually, one issue..." replies.

EXAMPLE — POST-PURCHASE SEQUENCE

Day 1: order confirmation with delivery estimate. Day 3: delivery confirmation plus usage tips. Day 10: check-in on results, AI answers questions. Day 25: running-low reminder with reorder link. Day 30: review request if happy, help if not. Every message two-way — the AI handles the replies.

Step 2: let the AI sell — carefully

The same conversations that answer questions can qualify leads (budget, timeline, service type), recommend the right service, book consultations directly and handle common objections. One rule: helpful, never pushy. The AI recommends what solves the customer's problem, not what pads the ticket — one pushy bot conversation can undo a year of brand-building.

Step 3: add channels one at a time

Master one channel before adding the next, keep answers consistent everywhere and expand toward wherever the client's volume actually is:

- SMS — usually first; highest open rates, pairs with missed-call text back
- Webchat — captures site traffic before it bounces
- Facebook Messenger and Instagram DMs — where younger customers already are
- Google Business Profile chat — high-intent searchers, straight from the listing
- WhatsApp — essential for international or service-trade audiences
- Voice AI — the after-hours phone answerer, and eventually the overflow line

Step 4: add languages without adding headcount

If the client serves a multilingual market, configure the AI to detect and respond in the customer's language, translate the core answers and have a native speaker sanity-check them. Compare that to $25–50K a year per multilingual hire — it's one of the easiest value stories in this whole guide, and a genuine differentiator for clients in the right neighborhoods.

Avoid the common pitfalls

Every one of these has burned an agency. Read them before your first launch, and again before your fifth — familiarity breeds shortcuts.

Pitfall 1: the bot sounds like a bot

Generic, obviously automated replies make the business feel cheap. Fix: personality from the voice worksheet, variation in common answers, conversational connectors ("I see," "great question") and contractions.

BEFORE "Your order has been confirmed. Delivery will be on May 15."

AFTER "You're all set! Your order's confirmed and should arrive May 15th. Anything else you'd like to know about it?"

Pitfall 2: the AI tries to handle everything

A bot that bluffs past its limits creates angrier customers than no bot at all. State its scope in the intro, give it a graceful bail-out line — "That one's beyond me, so I'm connecting you with a specialist who'll review our conversation and follow up shortly" — and enforce automatic escalation after repeated misses.

Pitfall 3: handoffs that make customers repeat themselves

Nothing says "we don't talk to each other" like asking for the order number twice. The thread is right there in Conversations — require the human to acknowledge it in their first message, and set response-time expectations at the moment of handoff.

Pitfall 4: outdated information

An AI confidently quoting last quarter's prices is worse than silence. Assign ownership and frequency for every information type:

INFORMATION	UPDATE FREQUENCY	OWNER
Availability and calendars	Live (calendar sync)	System
Pricing	Weekly check	Account manager
Promotions	At launch of each	Account manager
Policies	Quarterly	Client owner
Services offered	Monthly	Client owner

Pitfall 5: zero personalization

The CRM knows this contact's name, history and membership level — use it. "Welcome back, Taylor! I see you're on the premium plan — how can I help you get more out of it?" beats a cold greeting every time. Wire contact fields into the AI's context and add conditional responses for status and history.

AGENCY-SPECIFIC PITFALL

Deploying the snapshot and walking away. The clients who churn are the ones whose bot was never reviewed after launch. The monthly rhythm from chapter five isn't optional maintenance — it's the service.

09 — CONCLUSION

Ship it, then make it a product

Done right, AI customer service becomes an extension of your client's brand — representing them perfectly at 3am, in the tone they'd use themselves, booking the jobs they'd otherwise lose. And for you, it becomes something rarer: a deliverable that compounds. The benefits stack:

- **Consistency** — every customer gets the same right answer on every channel
- **Scalability** — capacity grows with demand, for the client and across your book of clients
- **Insight** — the conversation log is a running record of what customers actually want
- **Defensibility** — competitors can't match the service level without building what you built
- **Sanity** — owners disconnect at night; so does your team

Start simple, get the foundations right and improve from real conversations. The best implementations keep the split clean: AI for the routine, humans for empathy, judgment and exceptions.

Your four-week action plan

WEEK 1

Run the conversation audit on your pilot client (or your own agency inbox). Document the top 20 questions and complete the brand voice worksheet.

WEEK 2

Build in a template sub-account: train the knowledge base, write the model answers, wire the escalation workflows, connect calendars and channels.

WEEK 3

Closed beta in suggestive mode. Your team and the client's staff stress-test it; patch every miss. Capture the baseline metrics in writing.

WEEK 4

Go live on the first channel in auto-pilot. Start the weekly review rhythm, save the build as a snapshot and draft the offer for client number two.

ONGOING

Weekly conversation reviews. Monthly client reports. Quarterly capability expansion. Each deployment sharpens the snapshot — and raises what it's worth.

While your competitors are asleep, your clients' AI is answering questions, booking jobs and building their brand — and every one of those conversations is your agency's work, compounding. The only question is how many clients you'll have it running for by the end of the quarter.

READY TO BUILD IT?

Don’t build it alone. Join Nexus Hub

This guide gives you the playbook. Nexus Hub gives you everything else — the community of HighLevel agencies who've already deployed it, and the people to ask when you hit the weird edge case at step three.

- Weekly Q&A sessions with agency operators
- Exclusive snapshots and templates
- Premium partner deals
- A private community that answers fast

Join Nexus Hub

nexushub.club