MICHAEL
REIMER.AI

A FOUNDER GUIDE FROM MICHAEL REIMER

AI SYSTEMS IN PRODUCTION

Build, test and run agents on HighLevel

Michael Reimer

AI architect · Global keynote speaker · HighLevel Commander and Certified Admin · Founder of CRM Pros

BEFORE YOU START

THE BUILD IS THE EASY PART.

You did not have to hand over an email address to get this.
That was on purpose. I wanted you to have it with no strings.

I have spent about four years working full time in AI, building systems for my own companies and for the Founders I teach. This guide is about the part that comes after the demo.

A demo agent has to impress one person for five minutes. A production agent has to answer strangers every day, on its worst day, with nobody watching. Those are two different jobs, and most of the trouble I see comes from treating them as one.

My team works by a simple rule. Nothing goes live until it has been tested. Workflows are built in Draft and tested step by step before anyone publishes them. This guide shows you how to run that same discipline in your own account.

It has six chapters, and each one ends with a worksheet or a checklist. Pick one real system and fill them in as you go. Do that and you finish with a system that is designed, tested, signed off and owned.



Michael Reimer

AI architect · Global keynote speaker · HighLevel Commander and Certified Admin · Founder of CRM Pros

A plain disclosure. I am a sponsored HighLevel affiliate. When this guide points you to HighLevel, the link uses my partner link at no extra cost to you. Nothing in here is a promise of results. Your numbers depend on your business and on the work you put in.

INSIDE

SIX CHAPTERS. ONE DISCIPLINE.

01 **Production is different** Page 4

02 **Designing the agent** Page 8

03 **Building in HighLevel and AI Studio** Page 12

04 **Testing before launch** Page 16

05 **Running it every day** Page 20

06 **Scaling to many systems** Page 25

\+ **Where to go next** Page 28

The pages you fill in

Page 7 Production readiness score

Page 11 Agent design sheet

Page 15 Build checklist

Page 18 Test script table

Page 19 Launch sign off

Page 24 Weekly review sheet

Page 27 System register

Tip Print them once per system

THE IDEA IN ONE LINE

Nothing goes live until it has been tested, and nothing stays live without an owner.

01
PRODUCTION IS DIFFERENT.

A demo has to work once, for one person who wants it to work. A production system has to work every day, for people who do not care whether it works.

In a demo you choose the question. In production the customer chooses it, and they choose badly. They mumble. They change their mind halfway through. They ask about something you never thought to write down. They call at eleven at night.

A demo is judged by how good the best conversation is. A production system is judged by how bad the worst one is. So the work shifts. You spend less time making the agent sound clever and more time deciding what it does when it does not know.

A DEMO PROVES

The idea

It shows what an agent could do for the business. It runs on a script you control.

A TEST PROVES

The build

It shows the agent does what you designed, including when the caller goes off script.

PRODUCTION PROVES

The discipline

It shows the system keeps working after you stop watching it every hour.

None of this is an argument against demos. A good demo is how a Founder decides the system is worth building. The mistake is pointing the demo at real customers and calling it finished.

MICHAEL'S TAKE

The businesses that get burned by AI are the ones that let it guess. The ones that win set clear limits and a fast handoff.

SIDE BY SIDE

Demo agent versus production agent

Same platform, same tools and often the same prompt to start with. Here is what separates them.

	DEMO AGENT	PRODUCTION AGENT
Who talks to it	You, or a prospect you are guiding.	Strangers, on any day, in any mood.
What it knows	Sample facts, often written into the prompt.	The real business, held in a knowledge base that someone keeps current.
When it does not know	It improvises, and nobody minds.	It says so and hands off. It never guesses.
The handoff	Not needed. You are in the room.	A trigger, a route and a named person on the other end.
Testing	A few friendly runs.	A written test script, edge cases and a sign off.
After the conversation	Nothing happens.	A workflow confirms, updates the contact and creates a task for anything unresolved.
The rules	Ignored.	Consent, registration and disclosure handled before the first message.
The bill	Too small to notice.	Decided on purpose and checked every week.
The owner	Whoever built it.	One named person who reviews it every week.

Read down the right hand column. Very little of it is about the model or the prompt. It is about what surrounds the agent: what it knows, where it stops, who checks it and who owns it. That is the work this guide covers.

THE STANDARDS

Six standards a production system meets

I hold every system to the same six. If one is missing, the system is still a demo, however good it sounds.

1 **One job**
The agent has a single goal, usually a booking. An agent without a goal answers questions forever.

2 **Known answers**
The facts live in the knowledge base, in one place each. The instructions hold behavior only.

3 **A way out**
There is a handoff with written triggers, a route and a named person who receives it.

4 **Proof**
Someone other than the builder ran a written test script and signed off on the result.

5 **A watcher**
Real conversations are read every day in week one and every week after that.

6 **A record**
The system has a name that follows a pattern, a one page record and one owner.

Where each one is covered

The first three are design. You settle them on paper, and chapter two walks through them. The fourth is testing, in chapter four. The fifth is the daily and weekly review, in chapter five. The sixth is what lets you run more than one, in chapter six.

Chapter three sits in the middle because the build is the middle. It goes faster when the design is done first.

A fair warning. These standards slow down the first system. The design sheet, the test script and the sign off carry over to every system after it.

WORKSHEET

Production readiness score

Pick one agent you already have, or the one you are about to build. Score each standard. Zero means it does not exist. One means it exists in part. Two means it is written down and someone else could find it.

STANDARD	WHAT A TWO LOOKS LIKE	SCORE
One job	You can state the goal in one sentence.	0 1 2
Known answers	Prices, hours and policies sit in the knowledge base, not the prompt.	0 1 2
A way out	Handoff triggers are written and a named person receives them.	0 1 2
Proof	A filled in test script exists, with every line marked pass or fail.	0 1 2
A watcher	A weekly review is on the calendar with a name beside it.	0 1 2
A record	The system is named to a pattern and has a one page record.	0 1 2
	Total out of twelve	_____

0 to 5
It is a demo. Keep it away from real customers and start at chapter two.

6 to 9
It is close. Fix the lowest score first, then test again before launch.

10 to 12
It is ready for the launch sign off in chapter four.

My lowest score, and what I will do about it:

MICHAEL'S TAKE

Score it honestly. The number is for you, and a low score today is cheaper than a bad call next month.

02
DESIGNING THE AGENT.

An agent that fails in production was usually never designed. It was prompted. Design is the set of decisions you make on paper before you open the builder.

Start with one goal

Give the agent one outcome: book the appointment, qualify the lead or route the call. Everything else it does is in service of that. When a new request comes up later, the goal tells you whether it belongs in this agent or in another one.

Decide what is safe to hand over

AI agents are very good at pattern work and should never be given judgment work. For anything new, I ask three questions. A task that passes all three is a good fit. A task that fails any one of them gets a handoff.

QUESTION ONE

Is there a written answer?

If the right response exists in your knowledge base or your policies, AI can deliver it.

QUESTION TWO

What if it is wrong?

A wrong appointment time is an inconvenience. Wrong medical, legal or financial guidance can hurt someone.

QUESTION THREE

Would the customer mind?

Most people are happy to book with an AI at night. Few want one when they are upset.

AI HANDLES WELL

Hours, services and policy questions. Booking, rescheduling and confirming. Qualifying questions. Summaries of long conversations for your team.

PEOPLE KEEP

Legal, medical, tax or financial advice. Emergencies. Negotiation and exact quotes on complex work. Guarantees, refunds and exceptions. Upset customers.

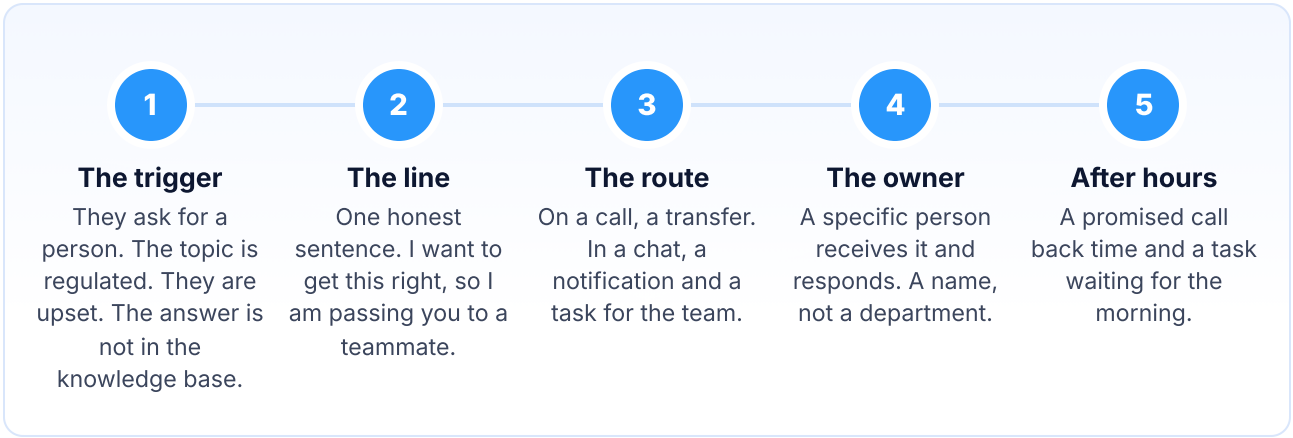
Guardrails are written, not assumed

A guardrail is a sentence in the instructions that tells the agent what it must not do. If it is not written, it does not exist. These are the ones I put in every agent.

- **Say what you are.** The agent says it is an AI assistant for the business in its first line.
- **Do not guess.** When the answer is not in the knowledge base, say so and hand off.
- **Stay out of advice.** No legal, medical, tax or financial guidance. Collect the details and pass them on.
- **Promise nothing.** No exact quotes on complex work, no guarantees and no exceptions to policy.
- **Use the approved wording.** Write the disclosure and the opt out wording into the instructions, so the agent never has to make it up.

The handoff has five parts

A handoff is more than a line in the prompt. It needs a trigger, a route and a person on the other end. Design all five parts, then test each one yourself.



MICHAEL'S TAKE

An unowned handoff is a dropped customer.
Name the person before you write the prompt.

Facts in one place, behavior in another

Mixing the two is the most common reason agents drift. When a price changes, you want to update one table and have both your text agent and your voice agent answer correctly.

THE KNOWLEDGE BASE HOLDS FACTS

Hours, prices, services, service areas, policies, staff roles and directions. In HighLevel you build it from your website through the web crawler, plus documents, question and answer pairs, and tables for structured data such as prices.

THE INSTRUCTIONS HOLD BEHAVIOR

Tone, the goal of the conversation, the order of the questions, what never to promise and when to hand off. If you find a price in the prompt, move it.

Trained web pages can refresh automatically. Build the knowledge base from the questions customers really ask, pulled from recent calls, texts and emails.

The prompt has six sections

SECTION	WHAT TO WRITE
Identity	Name, business, role, and that it is an AI assistant.
Goal	The one outcome: book, qualify or route.
Flow	The questions in order, asked one at a time.
Rules	Tone, length of replies, topics to avoid and what never to promise.
Handoff	The exact situations that trigger a transfer, and what to say first.
Close	Confirm the outcome, mention the text confirmation and thank them.

For voice, write for the ear. Keep replies to one or two short sentences. Ask one question per turn, then wait. Never read a list longer than two options. Tell the agent how to recover when it mishears. My team keeps HighLevel voice prompts under 4,000 characters. Anything you put inside quotation marks is read word for word, so quote with care.

WORKSHEET

Agent design sheet

One sheet per agent. Fill it in before you open the builder. If a row is hard to fill, that is the part of the design that is not finished.

System name	
Channels Phone, text, chat, social	
The one goal	
Questions, in order One per turn	
It must never	
Handoff triggers	
The handoff line	
Route and owner Transfer, or notify and task	
After hours plan	
Knowledge sources	☐ Website crawl ☐ Documents ☐ FAQs ☐ Tables
Mode at launch Text agents	☐ Suggestive, a person approves each reply ☐ Autopilot
After the conversation What the workflow does	

SYSTEM OWNER BUILDER TESTER

03

BUILDING IN HIGHLEVEL AND AI STUDIO.

Four pieces do the work. Each one has a narrow job. Most build problems come from asking one piece to do the job of another.

THE PIECE	WHAT IT DOES	STARTED BY
Conversation AI	Replies to contacts across SMS, email, web chat, WhatsApp, Facebook and Instagram. Answers from what it has been trained on, qualifies and books.	An incoming message
Voice AI	Answers inbound calls and places outbound calls. Answers from the knowledge base, books on your calendars, transfers calls and triggers workflows after each one.	A call, or an outbound workflow
Workflows	The automation around the agents: confirmations, tags, tasks and follow up. AI steps can sort a message by intent, summarize, translate or make a yes or no decision.	A trigger you choose
AI Studio	Builds the front door: websites, landing pages, multi step surveys, custom booking flows and branded forms, connected to your HighLevel form or calendar.	A prompt you write

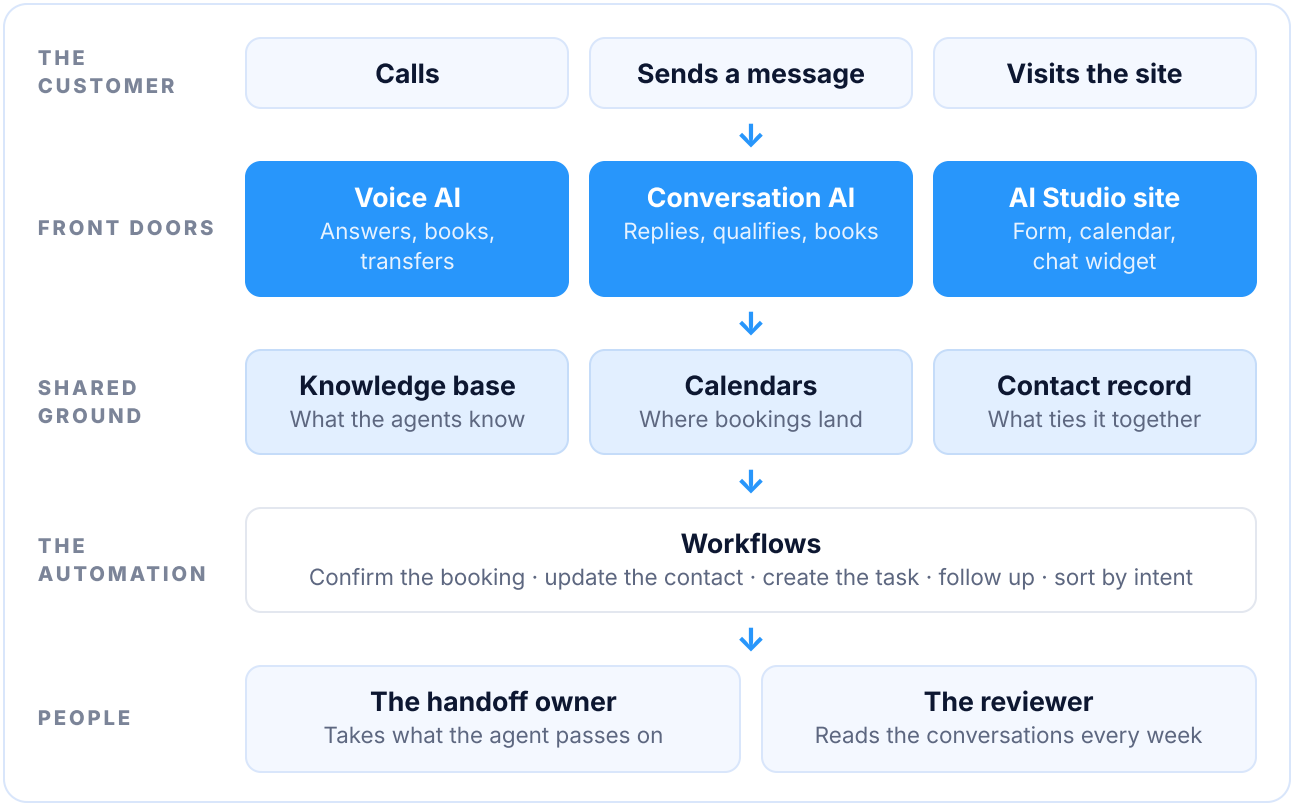
The agents talk to customers. Workflows do what happens before and after the conversation. AI Studio builds the page where the conversation starts. The contact record ties all of it together, which is the advantage of building on one platform.

MICHAEL'S TAKE

Use the simplest piece that does the job. A fixed rule beats an AI step when the rule is clear.

THE SYSTEM ON ONE PAGE

How the pieces connect



Four connections worth knowing

- **One knowledge base, two agents.** Conversation AI and Voice AI answer from the same knowledge base, so a fact is corrected once.
- **Workflows can steer the text agent.** A workflow can switch a Conversation AI bot on, off or between modes for a specific contact. A complaint tag can move one conversation to a person while the rest stay on autopilot.
- **Outbound calls start in a workflow.** Outbound Voice AI calls run from a workflow action, so the workflow decides who is called and when.
- **The site can hand chat to the agent.** An AI Studio site can connect its live chat widget to a Conversation AI agent from the site settings. The agent is still created and trained in Conversation AI.

The order to build in

Do not switch everything on in a weekend. Each agent needs a week of attention after launch, and you only have so much attention.

1. **The knowledge base.** Text and voice agents answer from it, so it comes before either.
2. **One customer facing agent.** Pick the channel where leads wait longest.
3. **The workflow around it.** Confirmation, contact update and a task for anything unresolved. Built in Draft.
4. **The other channel.** Reuse the rules and the knowledge you proved on the first.
5. **AI steps inside workflows.** They pay off most when the customer facing side is already steady.

Three settings that decide launch day

CONVERSATION AI

The mode

Off, suggestive or autopilot. In suggestive mode the agent drafts and a person sends. Start there, review for about a week, then switch to autopilot.

VOICE AI

The routing

The agent is half the setup. The other half is when a call reaches it. Pick one of the three patterns below and test it from your own phone.

WORKFLOWS

The status

Draft or Published. A workflow in Draft does not run for real contacts. Keep it there until the test script passes.

CALL ROUTING PATTERN	HOW IT IS SET
AI answers every call	Give the agent a number that is not your default subaccount number, with no call forwarding on it.
Team first, AI as backup	Forward to your team and set the inbound timeout, so the AI picks up before voicemail does.
AI after hours only	Set custom working hours on the agent in the Deploy tab.

Voice AI works on HighLevel phone system numbers or Twilio numbers. HighLevel moves its menus often. If a label has moved, search the setting name in the account search bar.

CHECKLIST

Build checklist

Work through it in order. Nothing on this page makes the system live. It makes the system ready to test.

KNOWLEDGE BASE

- [ ] Real customer questions collected from calls, texts and emails.
- [ ] Website added through the web crawler.
- [ ] Question and answer pairs written for the most common questions.
- [ ] Prices, service areas and schedules in tables.
- [ ] Old promotions and outdated policies removed.

CONVERSATION AI

- [ ] Channels chosen, starting with the busiest one.
- [ ] Mode set to suggestive.
- [ ] Brand voice and one goal set.
- [ ] Instructions cover what to avoid and when to hand off.
- [ ] A2P registration complete before any text to US mobile numbers.

VOICE AI

- [ ] Greeting says it is an AI assistant.
- [ ] Knowledge base attached.
- [ ] Appointment booking and call transfer added as actions.
- [ ] Number assigned and working hours set in the Deploy tab.
- [ ] Routing chosen in the call forwarding settings of the number.

WORKFLOWS

- [ ] Built and left in Draft.
- [ ] One job per workflow, with a clear name.
- [ ] Trigger filters name the exact form, calendar or tag.
- [ ] A branch exists for anything an AI step cannot sort.
- [ ] No workflow can trigger itself in a loop.

AI STUDIO PAGES, IF THE SYSTEM HAS ONE

- [ ] The HighLevel form or calendar is connected.
- [ ] Every fact, price and promise on the page was reviewed by a person.
- [ ] The mobile preview was checked.
- [ ] The chat widget points at the right agent.

Then stop. When every box is ticked, the builder hands the system to the tester. The workflows stay in Draft and the text agent stays in suggestive mode until the test script on page 18 passes.

04
TESTING BEFORE LAUNCH.

Nothing goes live until it has been tested. There is no exception for a small change, a rushed client or a build that looks right.

On my team building and testing are separate jobs. The builder leaves the workflows in Draft. The tester does not build anything. They push a contact through, one step at a time, and record what was expected and what happened. Every line gets a pass or a fail. The filled in sheet is the report.

Test in three layers

1	**Each piece**	One trigger, one action, one answer. Does this single step do what it should?
2	**The whole path**	One fresh contact, start to finish. Message, booking, confirmation, task. Push the contact past Wait steps by hand instead of waiting.
3	**The difficult customer**	The edge cases on the next page. This layer is where a demo and a production system part ways.

The tools HighLevel gives you

- **Workflows.** The Test Workflow feature in HighLevel Labs tests a single trigger or action with real events from the account, then runs the whole workflow step by step with a result for each step. These tests run the real steps, so messages can send and records can change. Use a test contact that belongs to you.
- **Voice AI.** Run web test calls from the agent screen, then call the number from your own phone at different times of day. Check the call log, transcript and summary for each test.
- **Conversation AI.** Message the business from a phone that is not an existing contact, on a channel the agent is set to answer.

EDGE CASES

Nine customers to meet before launch

You cannot judge an agent by reading its prompt. You have to be a difficult customer. Play each of these yourself, on the phone and in text, and write down what happened.

01
The simple booking
Wants an appointment and answers every question.
Pass: booked on the right calendar with the right details.

02
The price shopper
Asks what it costs before anything else.
Pass: answers from the knowledge base and never invents a number.

03
The urgent one
Says the words that mean it cannot wait.
Pass: says it is connecting them and the transfer happens.

04
The one who wants a person
Asks for a human in the first ten seconds.
Pass: hands off without arguing.

05
The interrupter
Talks over the agent and changes their mind halfway.
Pass: stops, listens and picks up the new request.

06
The mumbler
Gives an unclear name or address.
Pass: asks again instead of guessing.

07
The wanderer
Asks something unrelated, then comes back.
Pass: answers or declines, then returns to the goal.

08
The pauser
Stops in the middle of a sentence to think.
Pass: waits instead of cutting in.

09
The late caller
Makes contact outside working hours.
Pass: the after hours plan runs as designed.

Rules for a clean test

- **Use a fresh contact.** Old test contacts are often blocked by the re entry settings of a workflow, and a contact with do not disturb set will not get a reply.
- **Change one thing at a time.** If you change three settings and it improves, you will not know which one did it.
- **Compare versions.** Voice AI supports versioning. Make the change in a new version and test it against the current one before you publish.
- **Retest after every change.** A change to routing, numbers, hours, channels or the prompt sends the system back through the script.

WORKSHEET

Test script table

The tester fills this in, not the builder. Write what really happened, even when it passed. Add your own lines at the bottom for anything special to this system.

#	THE TEST	EXPECTED RESULT	WHAT HAPPENED	PASS	FAIL
1	A new contact messages or calls.	The agent answers in the right mode.		☐	☐
2	Ask a question the knowledge base covers.	A correct, short answer.		☐	☐
3	Ask a question it does not cover.	It says so and hands off.		☐	☐
4	Book an appointment.	Right calendar, right details.		☐	☐
5	Ask for a person.	The transfer or task reaches the owner.		☐	☐
6	Use an urgent trigger.	Immediate handoff.		☐	☐
7	Make contact after hours.	The after hours plan runs.		☐	☐
8	Enter a fresh contact into the workflow.	Each step fires, in order.		☐	☐
9	Push the contact past each Wait step.	The next step fires.		☐	☐
10	Finish a conversation.	Confirmation, contact update and task.		☐	☐
11	Message a contact who opted out.	Nothing is sent.		☐	☐
12	_____	_____		☐	☐
13	_____	_____		☐	☐

SYSTEM TESTED _____ TESTER _____ DATE _____

WORKSHEET

Launch sign off

Every box ticked, three names at the bottom, then publish. Not before.

THE AGENT

- [ ] Every line of the test script is marked pass.
- [ ] All nine edge cases were played by a person.
- [ ] The agent says it is an AI assistant in its first line.
- [ ] The agent hands off when the answer is not in the knowledge base.

THE AUTOMATION

- [ ] Each workflow passed step by step in Draft.
- [ ] Workflows are now Published and saved.
- [ ] Re entry settings were checked.
- [ ] The workflow that runs after the conversation fired in testing.

THE RULES

- [ ] Consent is recorded for everyone the system will call or text.
- [ ] A2P registration is complete for texting US mobile numbers.
- [ ] Callers are told the call may be recorded, where that applies.
- [ ] Opt outs and do not disturb are honored on every channel.

THE PEOPLE AND THE BILL

- [ ] The handoff owner received a test handoff and responded.
- [ ] Fifteen minutes a day is on the calendar for week one.
- [ ] The weekly review has a day, a time and a name.
- [ ] The AI plan and the Extra Usage setting were chosen on purpose.

This is a practical checklist, not legal advice. Rules on consent and recording differ by state and country. Have a qualified attorney review your consent language, scripts and recording notice before launch.

MICHAEL'S TAKE

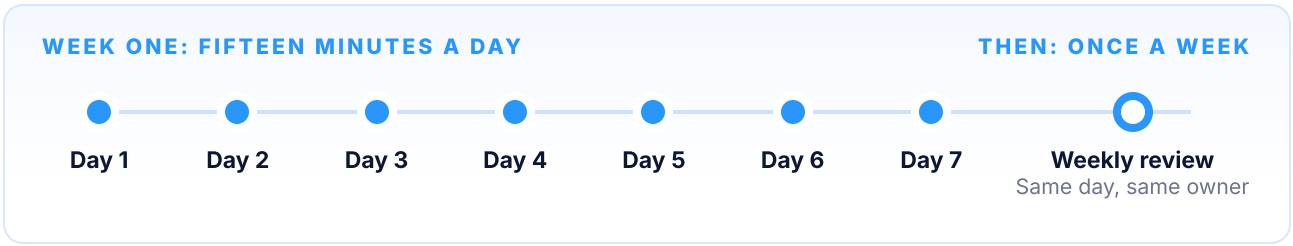
If a box cannot be ticked, the launch waits. A late launch costs a day. A bad one costs a customer.

BUILDER	TESTER	SYSTEM OWNER
GO LIVE DATE	FIRST WEEKLY REVIEW	WEEK ONE REVIEWER

05
RUNNING IT EVERY DAY.

The agent you launch is a first draft. The one you have a week later, after reading real conversations, is the one that earns its keep.

Launch day is not the finish. It is the first day you get real data. Set aside fifteen minutes a day for the first week. After that, a weekly review is enough.



What the fifteen minutes cover

1. **Read every transcript.** Look for the moment a customer got confused or repeated a question.
2. **Fix knowledge first.** If the agent did not know an answer, add it to the knowledge base. Do not patch facts into the instructions.
3. **Then fix behavior.** If the agent knew the answer and handled it badly, add one rule to the instructions.
4. **Check the bookings.** Each appointment landed on the right calendar with the right details.
5. **Check the handoffs.** Everyone who asked for a person reached one.
6. **Confirm the workflow fired.** Each conversation triggered what should happen after it.

By the end of the week the list of fixes gets short. That is your sign to move to the weekly rhythm.

The weekly review

Give each agent a few numbers and look at them every week. Then read the conversations behind the numbers. The numbers tell you where to look. The transcripts tell you what to fix.

THE PIECE	WHAT TO COUNT EACH WEEK
Conversation AI	How quickly it replies. How many conversations end in the goal. How many are handed to a person.
Voice AI	Calls answered. Calls that ended in a booking. Calls transferred. Callers who hung up early.
Workflows	Runs that failed, and the step that failed. Contacts that entered and stalled.

Voice AI gives you call logs with transcripts and summaries, plus performance reports. For workflows, Workflow Maps collects execution errors across a subaccount and shows the failed step. At launch it was open to agency users and agency admins.



1. **Read the numbers.** Mark each one up, flat or down against last week.
2. **Read the conversations.** A handful from each group: booked, handed off, dropped.
3. **Fix one thing.** Knowledge first, behavior second.
4. **Save it as a new version.** Test it against the current one, then publish.
5. **Compare next week.** Did the number move?

MICHAEL'S TAKE

A handoff is not a failure. A conversation that went quiet after a weak answer is.

DRIFT

Signs of drift, and where to look

Systems rarely break all at once. They drift. A price changes and nobody updates the table. A tag gets renamed. Someone adds a fact to the prompt. Look at the knowledge before you touch the prompt.

WHAT YOU NOTICE	WHAT IT USUALLY MEANS	THE FIX
Callers hang up early	The greeting is too long or unclear.	Shorten it. Say who is answering and ask one question.
Transfers that should have been bookings	A gap in the knowledge base.	Add the missing answer in one place, then ask again.
A wrong or old answer	Two sources disagree, or one is out of date.	Remove the old one. Turn on auto refresh for pages that change.
The voice agent talks over people	Long instructions, stacked questions or noisy calls.	Listen to the recordings. One question per turn. Change one setting at a time.
The text agent goes silent	It is off, in suggestive mode with drafts waiting, or not enabled on that channel.	Check status, mode, channels and which agent is primary.
The voice agent does not pick up	Not published, number not assigned, outside its hours or routed elsewhere.	Read the call log. If no AI session started, the call never reached the agent.
A workflow did not run	Still in Draft, a filter mismatch, or re entry blocked.	Compare the contact timeline with the workflow history.
Branches stop matching after a tag rename	Conditions and trigger filters kept the old tag name.	Update them by hand after every rename.

The tag rename trap. My team found that renaming a tag does not update the values inside If Else conditions or trigger filters. An If Else condition stores the tag name as plain text, so it can point at a tag that no longer exists. Check them against the live tag list after every rename and every build.

USAGE

The costs to watch

A production system has a meter. Three of them. Know which one each part of your system runs on, so the bill is something you decided and not something you discover.

CONVERSATION AI

By tokens

On pay per use it is billed on input and output tokens for the model you chose. Input includes the history, your instructions and knowledge base content. Text and messaging charges are separate.

VOICE AI

By the minute

Usage is measured in AI agent minutes, with inbound, outbound and widget calls combined. Phone system charges apply to every call on top of that, whatever the plan.

WORKFLOW AI ACTIONS

$0.01 each

Decision maker, intent detection, summarize and translate cost a cent per run, whatever AI plan the subaccount has. Steps that call an external model are billed by tokens.

OPTION, PER SUBACCOUNT	CONVERSATION AI	VOICE AI	PRICE
Pay per use	By tokens	By the minute	No monthly fee
AI Employee Growth	1,000 responses a month	100 minutes a month	$50 a month
AI Employee Unlimited	Unlimited, fair use	Unlimited, fair use	$97 a month

Rates are from HighLevel's published AI pricing as checked in November 2026. They can change, so confirm them on the current pricing page at MichaelReimer.AI before you quote anyone.

Four ways to keep it predictable

- **Put busy subaccounts on a plan.** A steady flow of calls and chats is more predictable on a plan than on pay per use.
- **Decide Extra Usage on purpose.** When a plan allowance runs out, usage pauses by default. Turn on Extra Usage only where you chose to pay for more.
- **Pick the model on purpose.** Smaller models cost a fraction per token and handle most booking and common question conversations well.
- **Watch for loops.** A workflow that triggers itself can run an AI step far more often than you planned.

WORKSHEET

Weekly review sheet

SYSTEM _____ WEEK OF _____ REVIEWER _____

NUMBER	LAST WEEK	THIS WEEK	UP	FLAT	DOWN
Conversations or calls handled			☐	☐	☐
Ended in the goal			☐	☐	☐
Handed to a person			☐	☐	☐
Dropped or hung up early			☐	☐	☐
Workflow runs that failed			☐	☐	☐
AI and phone usage this week			☐	☐	☐

What the conversations showed

Conversations I read
Booked, handed off, dropped

A question it could not answer

A moment a customer got confused

The one fix this week

Where the fix goes ☐ Knowledge base ☐ Instructions ☐ Workflow ☐ Routing

Before it is published ☐ Saved as a new version ☐ Retested ☐ Record updated

Owner of the fix, and by when

06

SCALING TO MANY SYSTEMS.

One system you can hold in your head. Ten you cannot. What carries you from one to ten is not a better prompt. It is names, records and owners.

Name everything to one pattern

A name should tell a stranger what the thing is and where it belongs, without opening it. My team uses one pattern across every build. This is how a workflow name is put together.

WF	01.01	Category	Name
What kind of asset it is	Folder number, then its order in the folder	The job family, such as booking	What this one does, in plain words

ASSET	PATTERN	WHAT IT BUYS YOU
Folders	`01 | Name`	Folders sort in the order you work.
Workflows	`WF | 01.01 | Category | Name`	Anyone can find the next step.
Tags	`category | detail`	Related tags sit together in every list.
Pipelines	`PIPELINE | Area | Name`	The area is clear before you open it.
Funnels	`FUNNEL | Niche | Type`	The same build is easy to reuse per niche.

The exact pattern matters less than having only one. Write yours down and refuse anything that does not follow it. A fixed scheme lets you reuse the same workflows for the next client and swap only the client facing content.

Two small helpers. A tag can carry a color and a workflow can hold sticky notes. Color the tag by category. Use the note to say what the workflow does.

Document the system on one page

Long documentation does not get read, and it does not get updated. One page per system does. It holds the answers a new person needs on their first day.

THE ONE PAGE RECORD HOLDS

The goal. The channels. Where the facts live. The handoff triggers, route and owner. The workflows it depends on, by name. The date of the last test and the last review.

WHERE TO KEEP IT

Next to the design sheet, the filled in test script and the sign off. One folder per system, named to the same pattern as everything else.

Three roles, one owner

One person can hold more than one role in a small team. What you cannot do is let the builder approve their own work, or let a system run with no name beside it.

ROLE ONE

The builder

Builds from the design sheet. Leaves workflows in Draft. Does not sign off on their own build.

ROLE TWO

The tester

Builds nothing. Pushes a contact through, fills in the test script and reports pass or fail.

ROLE THREE

The owner

Signs the launch. Receives the handoffs or names who does. Runs the weekly review.

The change rule

A small change that skipped testing is how a working system breaks. So every change, however small, takes the same path.

1. **One change**
   Never three at once.
2. **In Draft or a new version**
   Not on the live copy.
3. **Retest**
   The lines of the script it touches.
4. **Publish and record**
   Update the one page record.

WORKSHEET

System register

One line per live system. If a system is not on this page, nobody is watching it. Review the whole register once a month and retire anything that no longer has an owner.

SYSTEM NAME	CHANNEL	OWNER	LIVE SINCE	LAST TEST	LAST REVIEW

Before you add the next system

- [ ] The last system has run for a full week of daily reviews.
- [ ] Its weekly review has happened at least once, with a fix shipped.
- [ ] Its one page record is current and someone else has read it.
- [ ] The new system has a name that follows the pattern, and an owner.
- [ ] The owner has room for one more weekly review.

MICHAEL'S TAKE

You are not limited by how many systems you can build. You are limited by how many you can review.

NEXT STEP

PICK ONE SYSTEM. TAKE IT ALL THE WAY.

A guide does not put a system into production. A filled in test script and three signatures do. Here is the order I would go in.

1. **Score what you have.** Page 7. It takes ten minutes and it tells you where to start.
2. **Fill in the design sheet.** Page 11. One agent, one goal, one named handoff owner.
3. **Build it in Draft and run the test script.** Pages 15 and 18. Someone other than the builder tests.
4. **Sign it off and put the reviews on the calendar.** Pages 19 and 24. Daily for a week, then weekly.

If you want help building it

I build my systems on HighLevel, and I teach Founders to do the same. When you start HighLevel through my link you get a 30 day trial, and the CRM Pros team, the weekly calls and the community come with it, at the same HighLevel price. You can read exactly how that works, and what the plans include, at **MichaelReimer.AI**. More Founder guides are free at **MichaelReimer.AI**.

About Michael. Michael Reimer is an AI architect, global keynote speaker, HighLevel Commander and Certified Admin, and the Founder of CRM Pros. He has worked full time in AI for about four years, building AI systems for his own companies and teaching Founders and consultants to build theirs.

Disclosure. Michael is a sponsored HighLevel affiliate. Links to HighLevel use his partner link at no extra cost to you. This guide is educational. It is not legal, tax or financial advice, and it is not a promise of results. HighLevel product names are used to describe the platform as of November 2026 and can change.

© 2026 CRM Pros LLC. MichaelReimer.AI