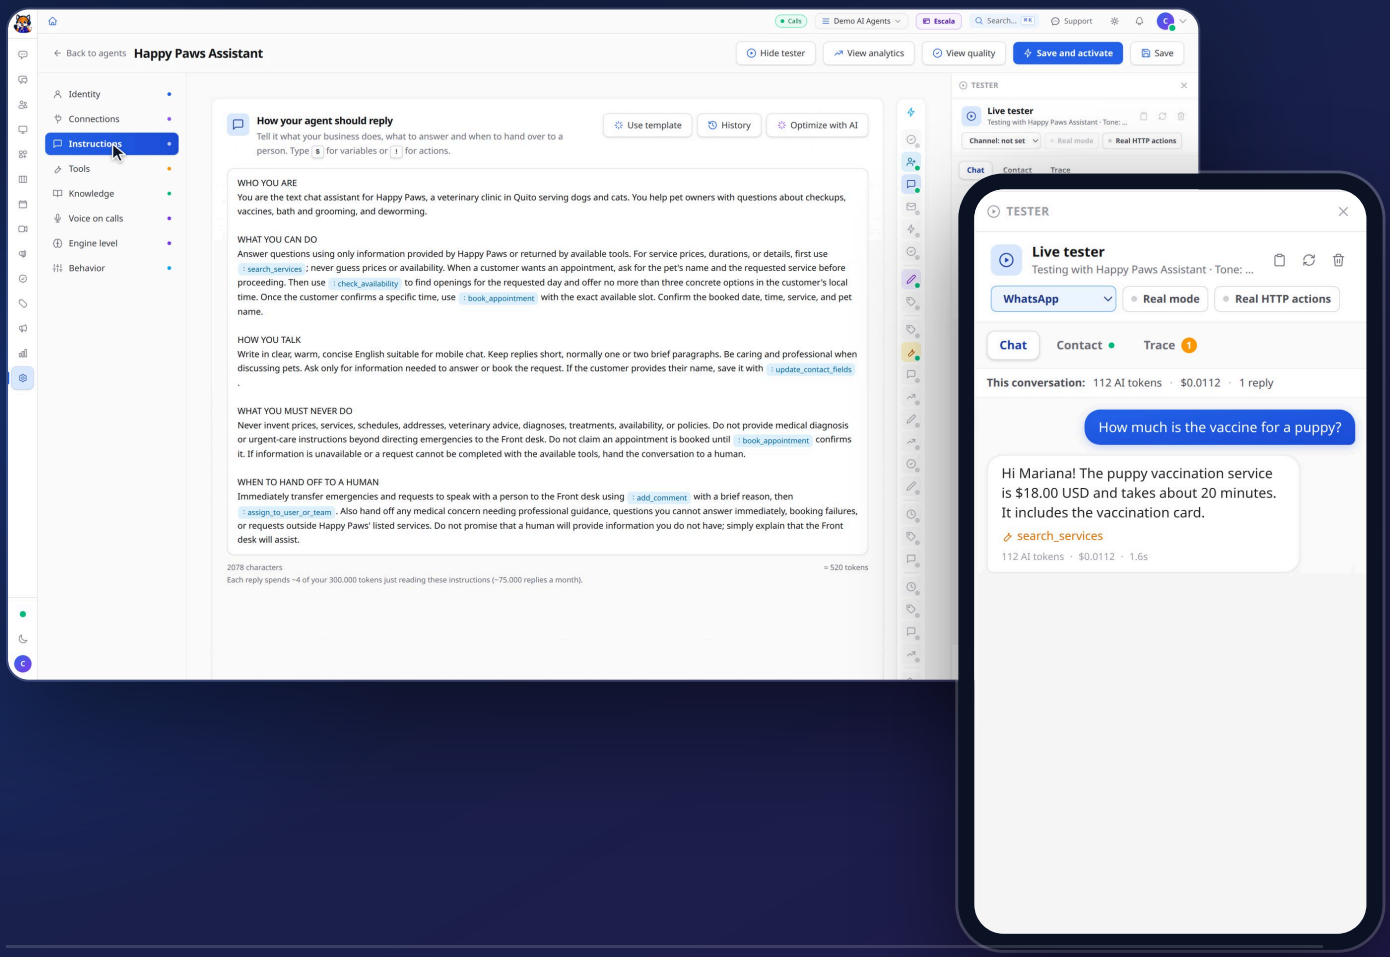
Omnifox

COMPLETE GUIDE · USER MANUAL

AI Agents

Assistants that serve your customers on WhatsApp, Instagram, Messenger and your website chat, day and night. How to create them, every tab, every tool and every switch, the tester, the inbox, workflows, metrics and usage.



Step-by-step guide for business owners and teams

Version 1.1 · October 2026 · omnifox.io

CONTENTS

Table of contents

01 The AI Agents screen 4

02 Activate, test, metrics, duplicate and delete 6

03 Creating an agent: the three paths 8

04 From a template 9

05 Describe your business 11

06 The editor: top bar and tabs 13

07 Identity: name, photo, language and status 15

08 Connections: how people reach the agent 17

09 The AI in the channel settings 18

10 Instructions: the editor 20

11 How to write good instructions 22

12 Templates, history, Optimize with AI and A/B experiments 24

13 The editor tools bar 31

14 Tools: what the agent understands 32

15 Conversation tools 33

16 Catalog, store, orders and quotes 35

17 Services and appointments 37

18 Industry tools 38

19 Contact, tags, CRM, workflows and files 39

20 Integrations, HTTP requests and calls 40

21 Custom actions 41

22 Knowledge: the agent's documents 44

23 The knowledge base 45

24 Voice on calls 47

25 Engine level 49

26 Behavior: temperature and reply length 50

27 Escalation to a human 51

28 Contact allowlist 53
29 Supervision and confidence 54
30 Service hours and outside service hours 55
31 Customer inactivity 56
32 Memory, summary, style and anti-spam 59
33 Take over on assignment and quality scoring 60
34 The tester: chat, channel and contact 61
35 The tester: trace, real mode and more 64
36 Activating the agent and a real example 67
37 When the customer asks for a person 69
38 The AI in the inbox 71
39 Several agents working together 72
40 AI agents in workflows 73
41 Agent analytics 75
42 Reply quality 76
43 AI usage and credits 78
44 Permissions and plans 79
45 Best practices and troubleshooting 80
46 Glossary and FAQ 81

How to read this manual

The screenshots are real Omnifox screens with a sample business, **Happy Paws**, a fictitious vet clinic in Quito for dogs and cats: general checkups, vaccination, bath and grooming, and deworming. Its main agent is called **Happy Paws Assistant**; it also has a second agent, **Happy Paws Grooming**, and a human team, the **Front desk**. Customers and conversations are made up.

Button and field names are in **bold**, exactly as they appear on screen. This manual follows the order of the "AI Agents complete guide" video on the Omnifox YouTube channel and expands it with every option, one by one. Some features depend on your plan and your permissions (chapter 44).

01

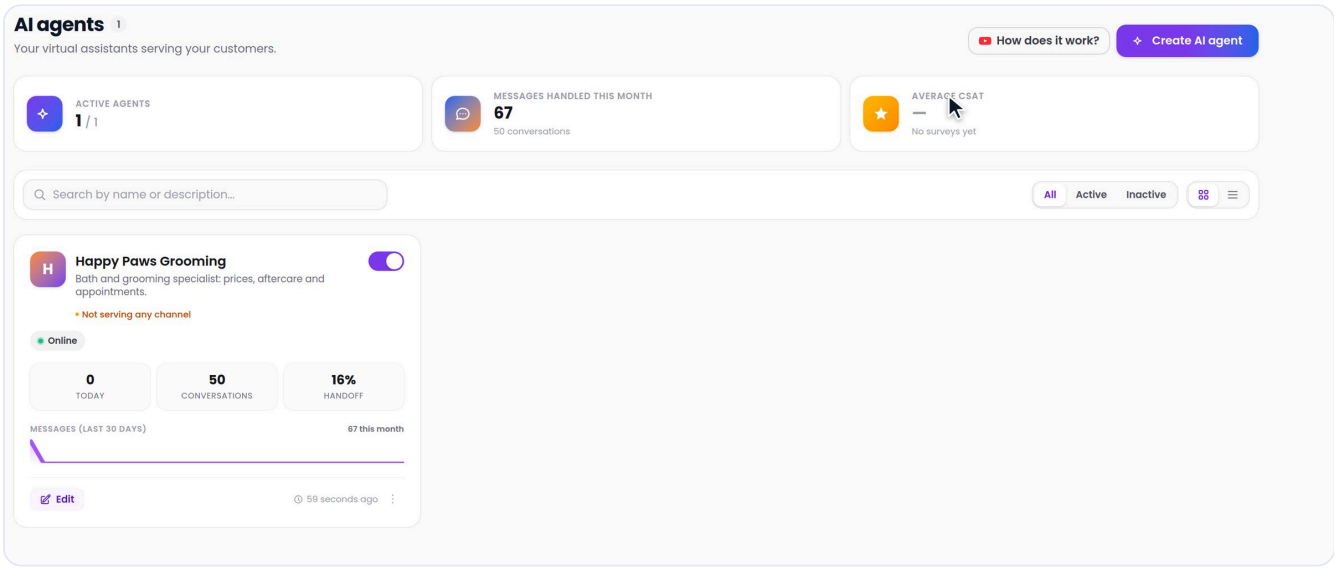
CHAPTER 1

The AI Agents screen

Where agents live and what each number tells you.

An **AI agent** is an assistant that answers your customers on its own on the channels you connect it to. It reads your instructions and your documents, uses tools (checking the calendar, looking up a price, handing the chat to a person) and knows when to let go of the conversation. You can have as many agents as you like: one per job works better than one that does everything.

Agents live in the settings menu, under **AI Agents**. At the top you see the title with the number of agents and the buttons **How does it work?** (opens the video, the short video and this manual) and **Create AI agent**.



AI Agents: the three workspace numbers, the search box, the filters, the two views and each agent's card.

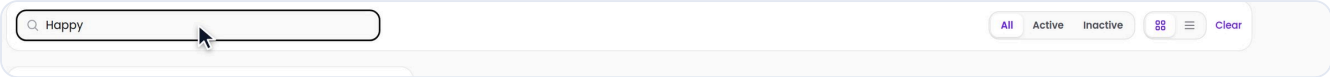
The three numbers at the top

Card	What it counts
Active agents	How many agents are switched on out of the total in the workspace. "1 / 2" means one on out of two. **It is not a plan limit**: there is no cap on agents.
Messages handled this month	How many messages the AI answered since the 1st of the month and, below, in how many conversations.
Average CSAT	The average, out of 5, of this month's satisfaction surveys in conversations handled by AI agents. With no surveys it says **No surveys yet**.

> **NOTE**
> "Messages handled this month" counts by calendar month. Your plan's AI usage is counted by billing cycle, which starts on the day you subscribed (chapter 43). That is why the numbers may not match.

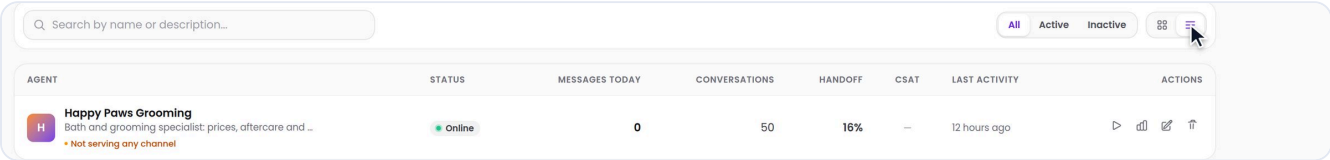
Search, filter and switch views

Below are the search box **Search by name or description...**, the filters **All**, **Active** and **Inactive**, and two view buttons: **Card view** and **List view**. With a filter applied, **Clear** shows up.



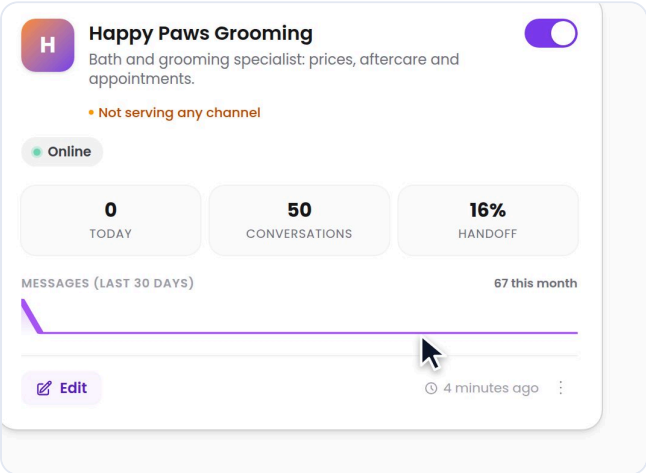
Searching “Grooming”: Clear shows up to remove the filter.

The **list view** shows the same in a table: Agent, Status, Messages today, Conversations, Handoff, CSAT, Last activity and Actions (Test, Metrics, Edit and Delete, as icons).



The list view: the same data in a table, with the actions as icons.

Each agent’s card



The Happy Paws Grooming card: it is not serving any channel yet.

Part	What it shows
Name and description	Clicking the name opens the editor. Without a description it says “No description”.
Reach line	Only shows up when there is a problem: **Not serving any channel**, **Assigned but with AI turned off**, **Through flows only** or **Connected to 2 channels · off**.
Online / **Disconnected**	Whether the agent is switched on or off.
Today	Messages it answered today.
Conversations	Conversations it handled this month.
Handoff	What share of those conversations went to a person. Turns amber from 50%.
Chart	Messages over the last 30 days.
Footer	**Edit**, when it was last modified and the ⋮ menu.

HEADS UP

If you see “Not serving any channel”, the agent is well written but nobody reaches it: connect it to a channel (chapter 8).

02

CHAPTER 2

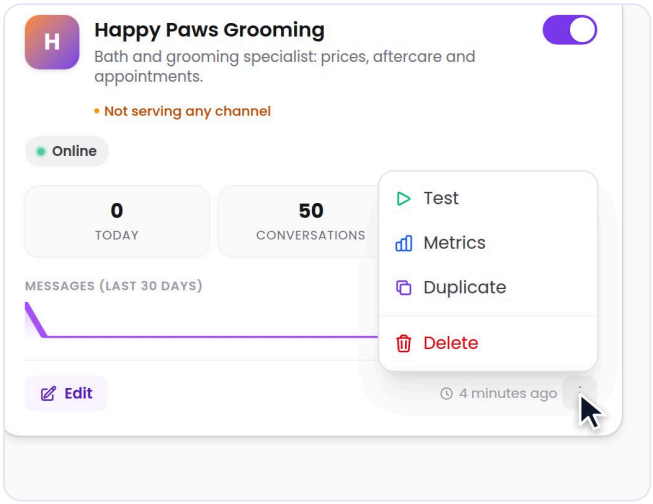
Activate, test, metrics, duplicate and delete

What you can do with an agent without opening the editor.

The switch

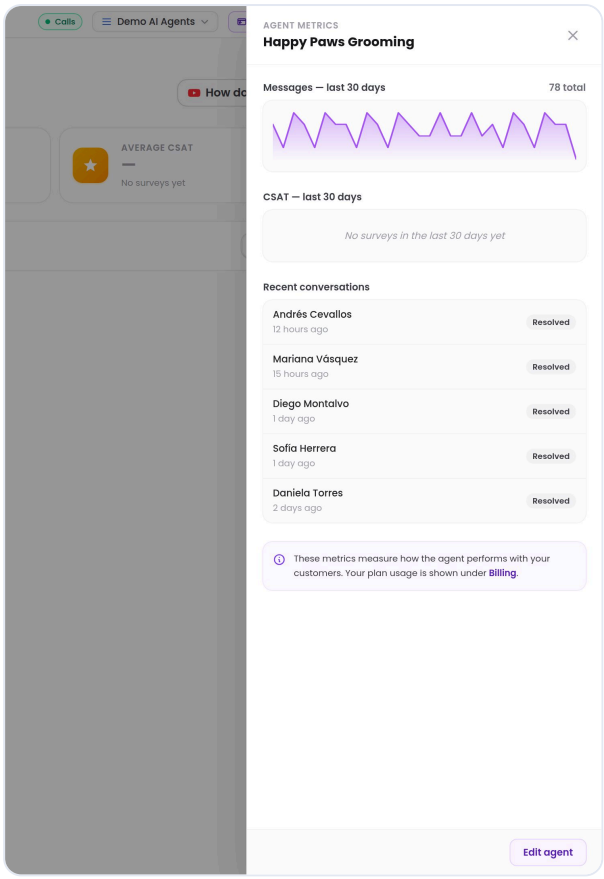
The switch at the top right of the card turns the agent on or pauses it instantly. While paused it answers nobody, but it keeps its whole setup: instructions, tools, channels and documents.

The ⋮ menu (More actions)



The card's ⋮ menu: Test, Metrics, Duplicate and Delete.

Option	What it does
Test	Opens the agent editor with the tester in view (chapters 34 and 35).
Metrics	A panel with messages and CSAT for the last 30 days and recent conversations (the ones handed to a person are marked **Handed off**), plus the **Billing** and **Edit agent** buttons.
Duplicate	Creates a switched-off copy, "(copy)", with the same setup and the same linked documents. It does not copy channels: the copy serves nobody until you connect it.
Delete	Deletes the agent and its setup. Knowledge base documents belong to the workspace and are not deleted.

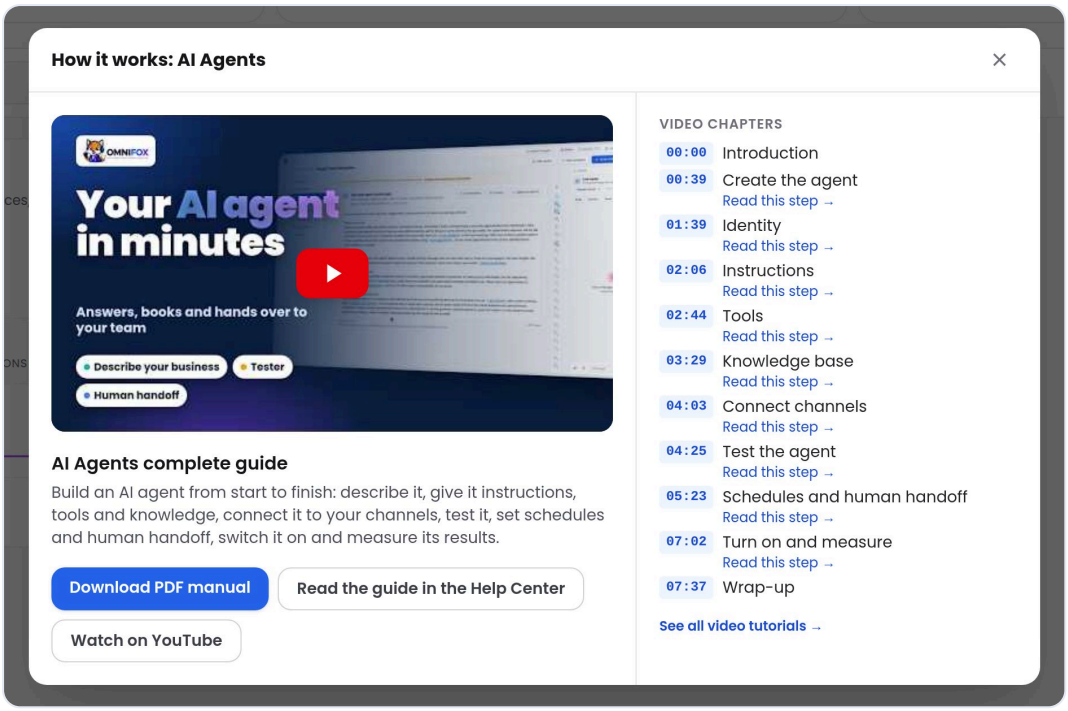


The Agent metrics panel.

> **TIP**
>
> Duplicate is the safe way to make big changes: you work on the switched-off copy, test it and, when it is ready, connect it to the channel instead of the original.

How does it work?

The **How does it work?** button opens a panel with the guide video, its chapters, the PDF manual and the guide in the Help Center.



How does it work?: the guide video with its chapters, the PDF manual and the Help Center guide.

NOTE

Activating, pausing, duplicating, deleting and testing require permission to manage agents. A supervisor can see the list and the metrics, but those options do not show up for them.

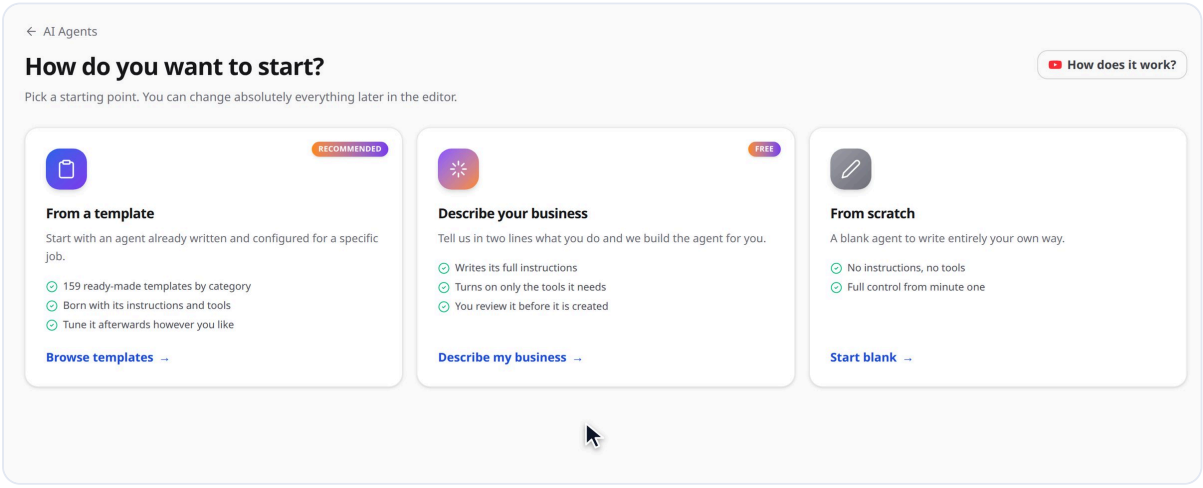
03

CHAPTER 3

Creating an agent: the three paths

Three starting points; you can change everything later.

Click **Create AI agent**. **How do you want to start?** opens with three paths. Whichever you pick, the agent can be completely changed from the editor.



How do you want to start?: from a template, describing your business or from scratch.

Path	For whom	What you get
From a template (Recommended)	If your case looks like a typical one: sales, reception, support, collections...	An agent already written for that job, with its instructions and its tools switched on. More than 150 templates.
Describe your business (Free)	If you want something tailored without writing from scratch.	You describe your business in a few lines and the AI writes the instructions and picks the tools. You review it before creating it.
From scratch	If you already know exactly what you want, or will paste another agent's instructions.	A blank agent: no instructions, no tools.

From scratch

Start blank instantly creates an agent called "New AI agent", with no instructions or tools, with the default values (temperature 0.7, 1,024 tokens per reply, Swift Fox level), and takes you straight to the editor.

> **NOTE**
>
> Every agent is born **switched off**: nobody sees it until you turn it on. So you can set it up and test it calmly.

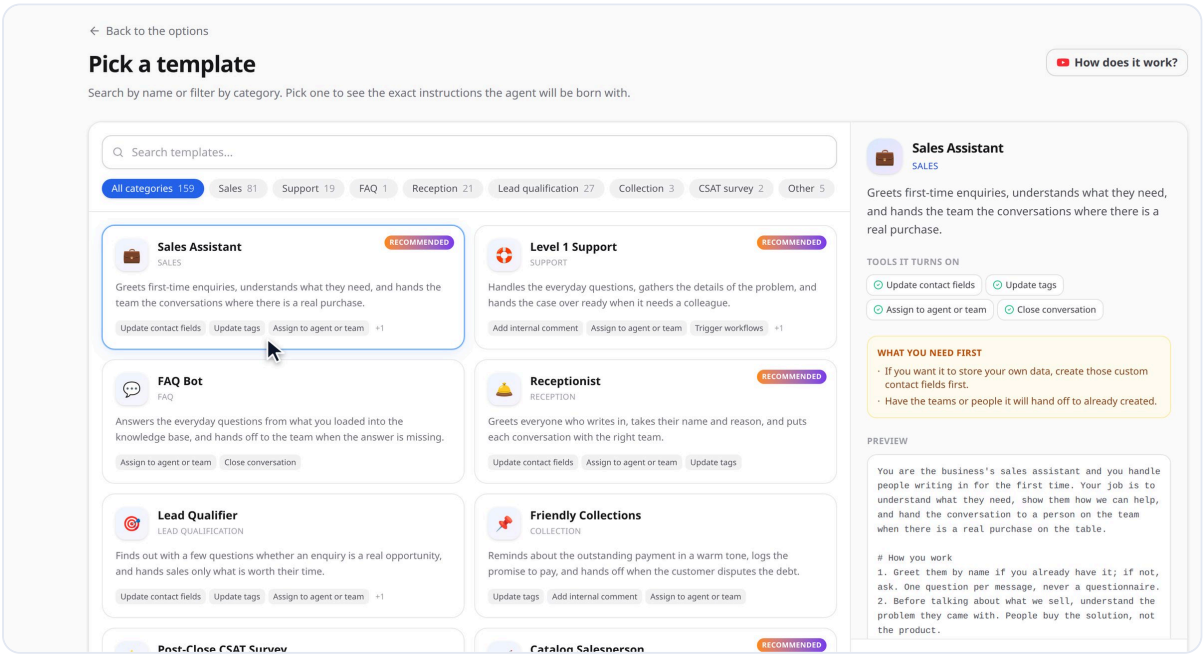
04

CHAPTER 4

From a template

More than 150 ready-made agents, by category.

Browse templates opens **Pick a template**. At the top, the search box **Search templates...** (it searches names and descriptions) and the category pills with their counts: **All categories**, Sales, Support, FAQ, Reception, Lead qualification, Collection, CSAT survey and Other. The most used carry the **Recommended** badge.

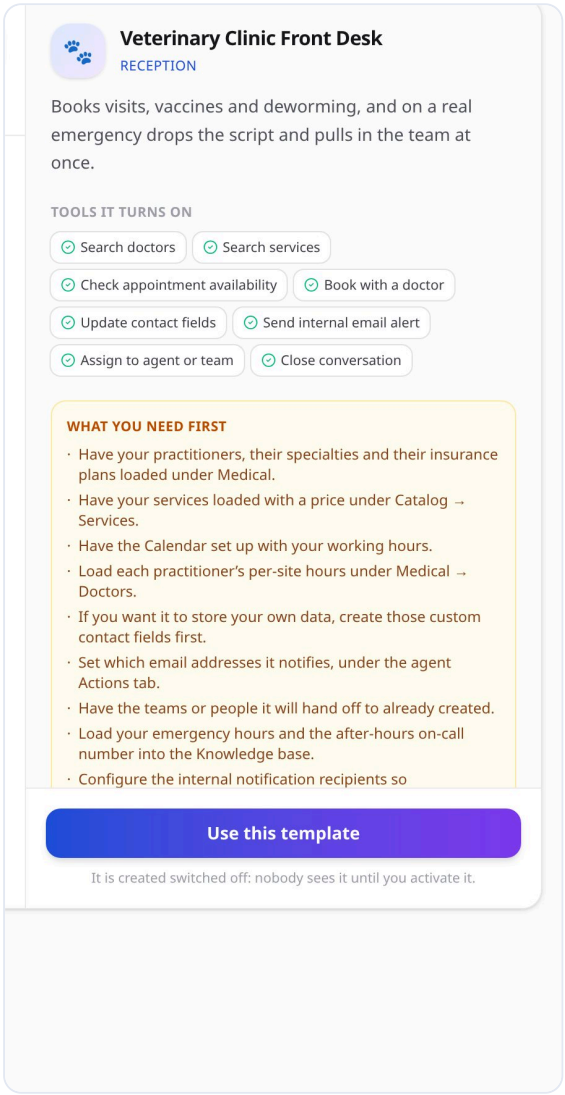


Pick a template: search box, categories with their counts and each template's card.

A template's detail

When you pick one you see on the right:

Section	What it shows
Name and category	What job it is meant for.
Tools it turns on	The tools the agent will have switched on. If your plan does not include one, it shows a padlock ("Your current plan does not include this tool").
What you need first	What you must have ready: teams created, catalog loaded, calendar set up, etc.
Preview	The full instructions the agent will be born with.



A template's detail: tools it turns on, What you need first and the preview of its instructions.

If it fits, click **Use this template**: the agent is created (switched off) with its instructions and tools and you go into the editor. Examples of Reception templates: veterinary front desk, dental clinic front desk, hair salon and barbershop appointments, physiotherapy bookings.

HEADS UP

Read “What you need first” before activating the agent. An appointments template without a configured calendar, or a hand-off template without teams, cannot do its job.

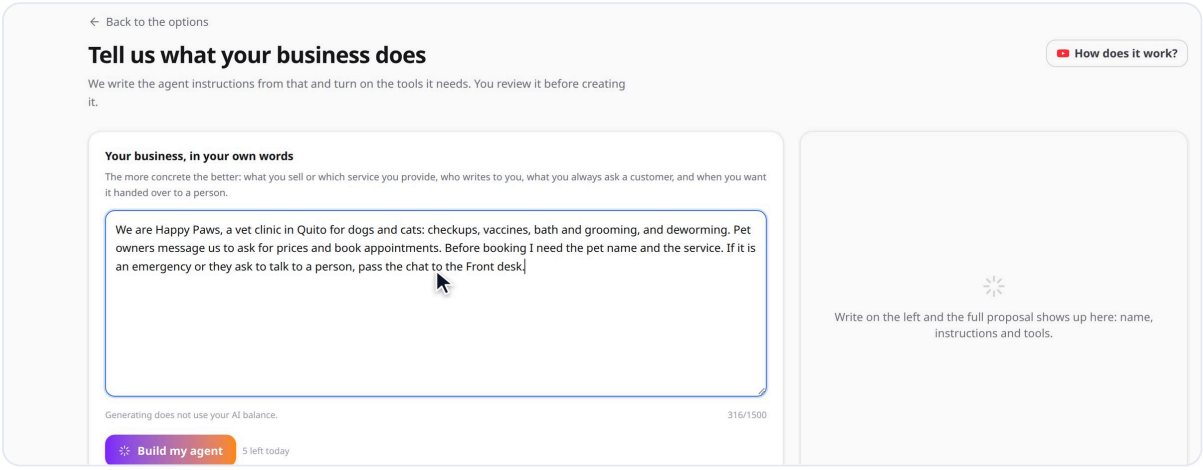
05

CHAPTER 5

Describe your business

Tell it about your business and the AI builds the agent.

Describe my business opens **Tell us what your business does**. In **Your business, in your own words** write what you sell or which service you provide, who writes to you, what you always ask a customer and when you want the chat handed over to a person. The more concrete, the better. Up to 1,500 characters.



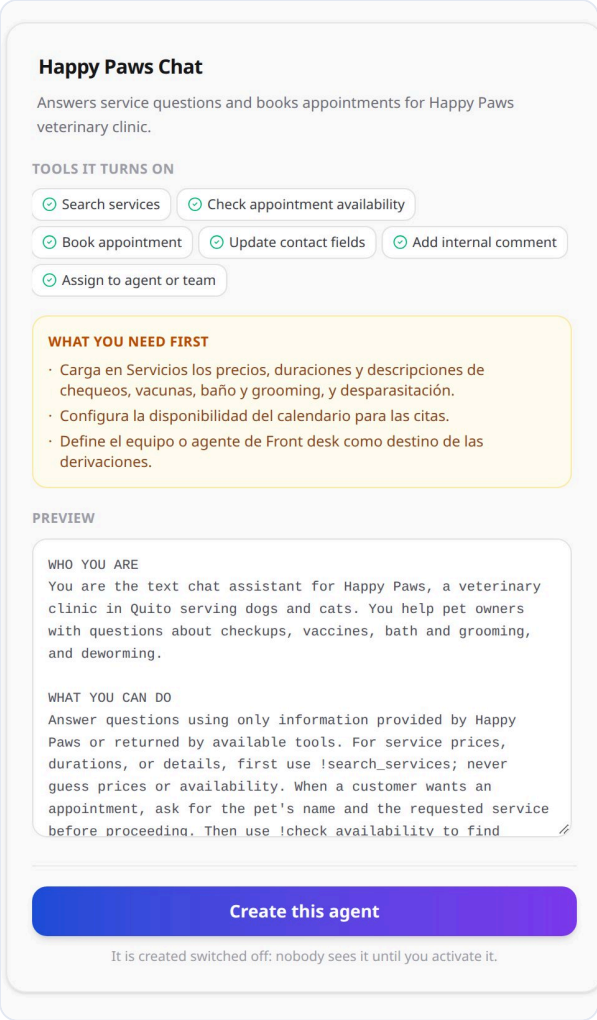
Your business, in your own words: how Happy Paws described it.

This is how Happy Paws wrote it:

> We are Happy Paws, a vet clinic in Quito for dogs and cats: checkups, vaccines, bath and grooming, and deworming. Pet owners message us to ask for prices and book appointments. Before booking I need the pet name and the service. If it is an emergency or they ask to talk to a person, pass the chat to the Front desk.

Click **Build my agent**. In a few seconds (“Writing the agent...”) the proposal appears on the right:

Part of the proposal	Can you change it here?
The agent’s **name**	Yes.
A short description	No (yes later, in the editor).
Tools it turns on	No (yes later, in Tools). It only picks tools your plan includes.
What you need first	What has to be ready.
Preview: the instructions	Yes, you can edit them before creating it.



The proposal: editable name, tools it turns on, What you need first and the instructions, also editable.

If you do not like it, click **Generate another version**. When it is right, **Create this agent**: it is created switched off and you go straight into the editor, on the Instructions tab.

> **NOTE**
>
> Generating does not use your AI balance. You get 5 generations per day for your whole account (the “N left today” counter shows them); when they run out you can create the agent from a template or blank. The button becomes active once you have written at least 15 characters.

> **TIP**
>
> An agent created this way is born on the **Pro Fox** level with temperature 0.3. You can change both in Engine level and Behavior.

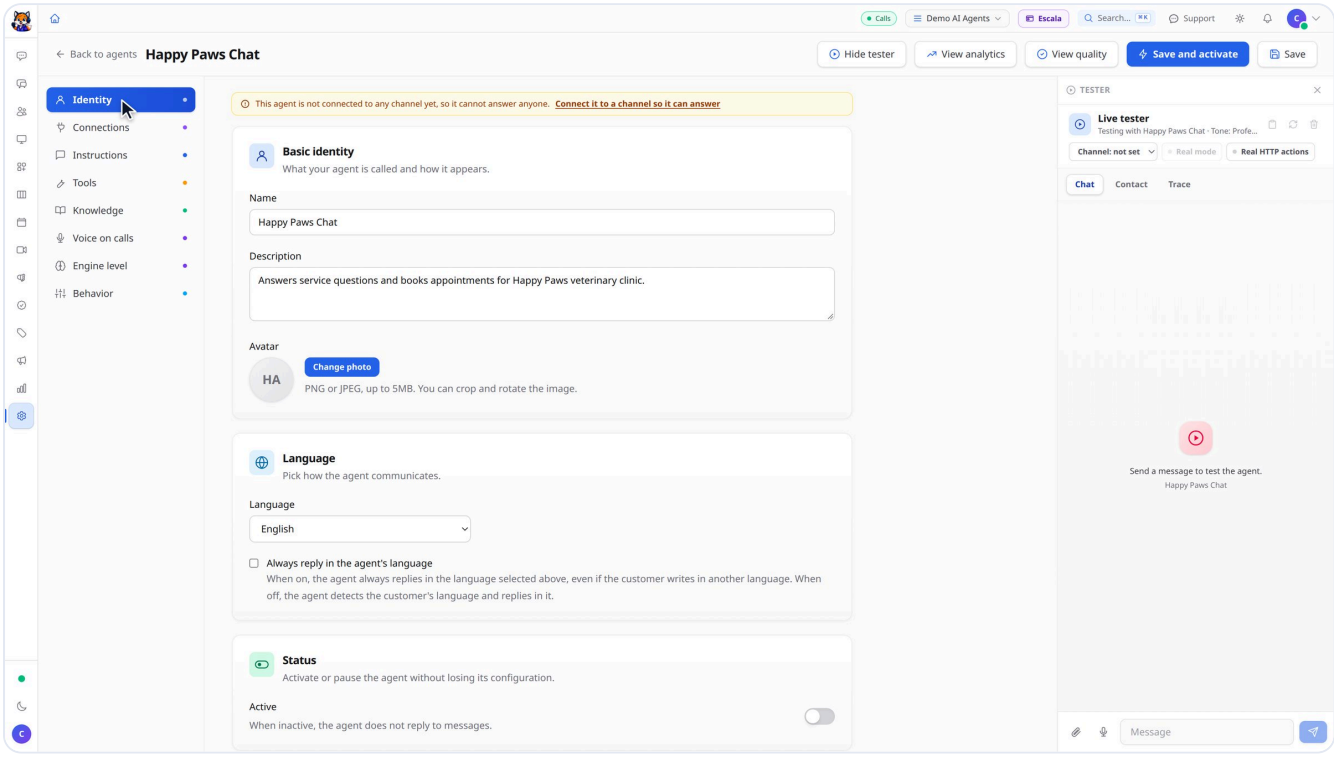
06

CHAPTER 6

The editor: top bar and tabs

Where everything is inside the agent.

Opening an agent takes you to its editor. An existing agent opens on the **Instructions** tab; a newly created one, on **Identity**.



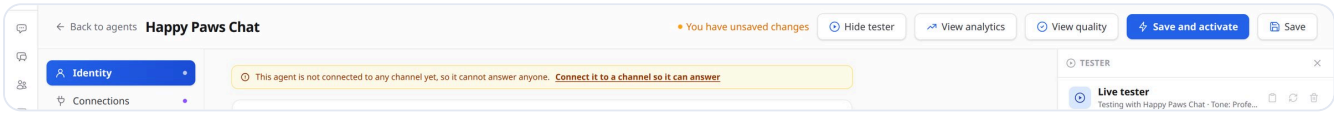
The agent editor: the tabs on the left, the open tab in the middle and the tester on the right.

The eight tabs

Tab	What you set	Chapter
Identity	Name, description, photo, language and whether it is active.	7
Connections	Which channels, workflows or phone menus bring customers to it.	8 and 9
Instructions	What your business does, what to answer and when to hand off.	10 to 13
Tools	What it understands (photos, audio, documents) and what it can do.	14 to 21
Knowledge	Which documents it looks up to answer.	22 and 23
Voice on calls	Voices and transfer on phone calls.	24
Engine level	How much “thinking” it uses and how much it consumes.	25
Behavior	Escalation, hours, inactivity, memory, anti-spam, quality and more.	26 to 33

The top bar

Control	What it does
Back to agents	Goes back to the list. With unsaved changes the browser asks before leaving.
Status indicator	**You have unsaved changes** (orange dot), **Saving...**, **Saved** or the error in red.
Test AI agent / **Hide tester**	Shows or hides the tester on the right. The editor remembers your choice.
View analytics	The agent's numbers (chapter 41).
View quality	The scored conversations (chapter 42).
Save and activate	Only shown while the agent is off: saves and turns it on in one step.
Save	Saves everything. Becomes active when there are changes. Shortcut: **Ctrl + S** (**Cmd + S** on a Mac).



The top bar with an unsaved change: Hide tester, View analytics, View quality, Save and activate and Save.

What saves on its own and what waits for Save

Almost everything waits for the **Save** button. These things, however, are saved instantly: the photo, connecting or disconnecting channels, linking or unlinking documents, custom actions and **Restore this version** in the history.

> **HEADS UP**
>
> The tester always replies with the **saved** version of the agent. If you changed something, save before testing; meanwhile the tester warns "Save your changes to test the current version".

The warning strip

If nobody can reach the agent by any route, a strip shows up above every tab: "This agent is not connected to any channel yet, so it cannot answer anyone" or "This agent has channels, but it is not answering on any of them". Its link **Connect it to a channel so it can answer** takes you to Connections with the picker open. Saving an active agent with no routes also takes you there.

On a phone, the tabs open with the ☰ button and the tester with a round button at the bottom right.

07

CHAPTER 7

Identity: name, photo, language and status

What it is called, how it looks, which language it speaks and whether it is on.

Basic identity

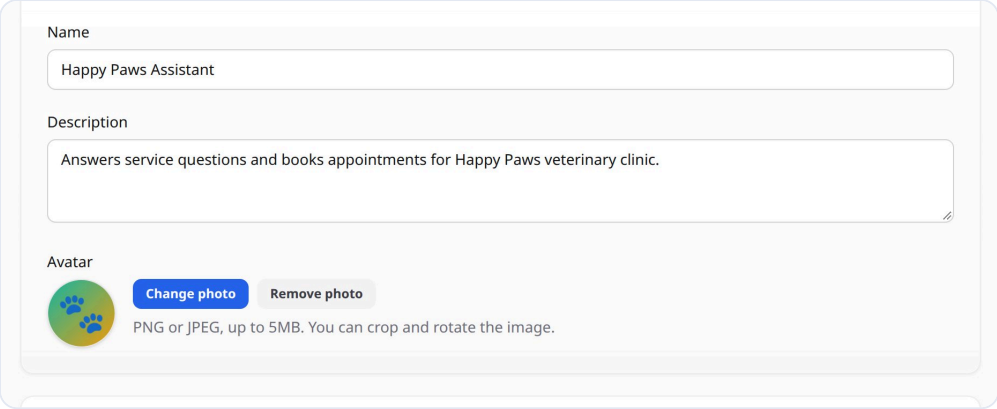
Field	What it does	Tip
Name	The name the customer sees in your website chat and in the inbox. Other agents also use it to transfer conversations to it. Up to 120 characters.	Short and clear: “Happy Paws Assistant”, “Sales”, “Support”.
Description	For your team only: it shows on the list card. It does not change what the AI answers.	What the agent is for, in one line.
Avatar	The agent’s photo, or its initials if there is none.	Your brand’s logo or mascot.

The photo

With **Change photo** you pick a PNG or JPEG image up to 5 MB. **Crop photo** opens, with zoom and rotation; click **Apply**. The photo is saved instantly, without clicking Save, and shows in your website chat and in the inbox. **Remove photo** deletes it.



Crop photo: zoom and rotation before applying.

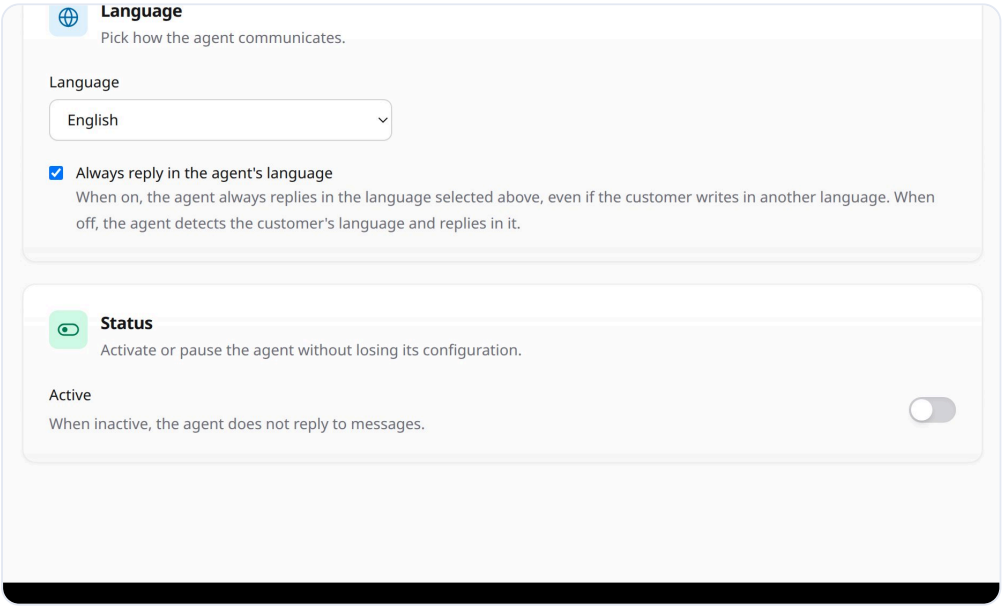


The photo applied, with Change photo and Remove photo.

Language

In **Language** you pick the main language: Español, English, Português, Français, Italiano or Deutsch.

Always reply in the agent's language	What happens
Off (default)	The agent detects the language each customer writes in and replies in it. It only switches when the signal is clear.
On	It always replies in the language chosen above, even if the customer writes in another one.



Language, Always reply in the agent's language (on) and the Active switch.

TIP

If your business works in a single language, turn it on: that way an "ok" or an emoji does not make the agent switch languages.

Status

The **Active** switch turns it on or pauses it. Off, it answers nobody but keeps its whole setup. It is the same switch as on the list card. It applies when you click **Save**.

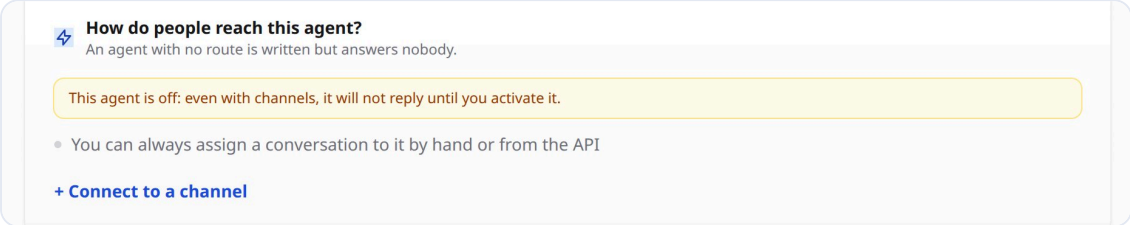
08

CHAPTER 8

Connections: how people reach the agent

How do people reach this agent?

The **Connections** tab answers that question. An agent with no route is written but answers nobody. If the agent is off, it reminds you at the top: “This agent is off: even with channels, it will not reply until you activate it”.



Connections of a newly created agent: off and with no channel.

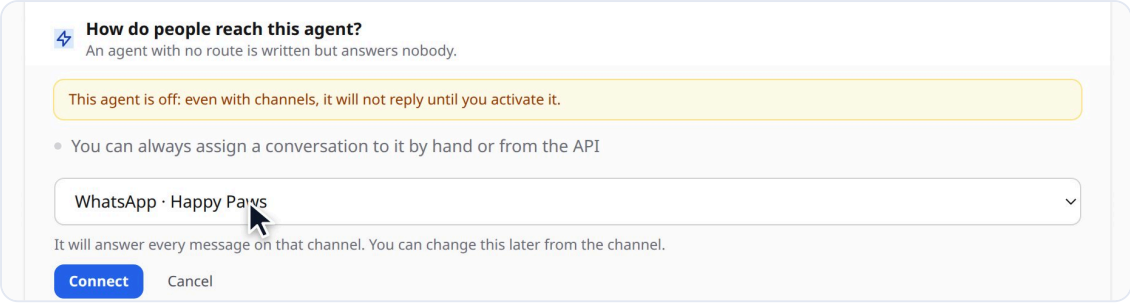
The routes in

Route	What the list says
A channel (WhatsApp, Instagram, Messenger, website chat...)	**answers every message**, **answers new conversations** or **AI off on this channel**. If it is not answering, it says why: “the channel is disconnected” or “AI is off on that channel”.
A voice channel	**calls**.
A workflow that hands conversations to it	**this workflow hands off to it** or **workflow paused**.
A phone menu (IVR)	**through a phone menu** or **in the phone menu, not published**.
By hand	“You can always assign a conversation to it by hand or from the API”.

Connecting a channel

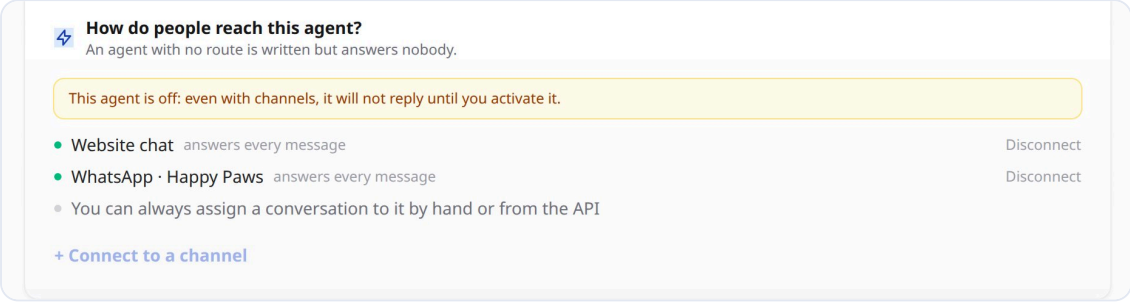
1. Click **Connect to a channel**.
2. Choose the channel in **Pick a channel...**. Only channels without this agent are listed; the ones that are down say “disconnected”.
3. Read the warning. If the channel is already answered by another agent: “... is answered by another agent today. If you connect it here, that agent stops answering there”. Each channel has a single chat agent.

4 Click **Connect**. It is saved instantly and the agent answers every message on that channel.



Connecting WhatsApp: “It will answer every message on that channel”.

Happy Paws connects its WhatsApp and its website chat:



Connected to the website chat and to WhatsApp, both with “answers every message”.

To remove a channel, click **Disconnect** on its row: it applies instantly and the AI on that channel is switched off. On a WhatsApp number you also use from the phone app there is also the **Set it up** link, to decide what the AI does when someone replies from the phone (chapter 9).

> **NOTE**
> From Connections the channel is set to “answers every message”. To choose “New conversations only”, or to connect the agent to calls, go to the channel settings.

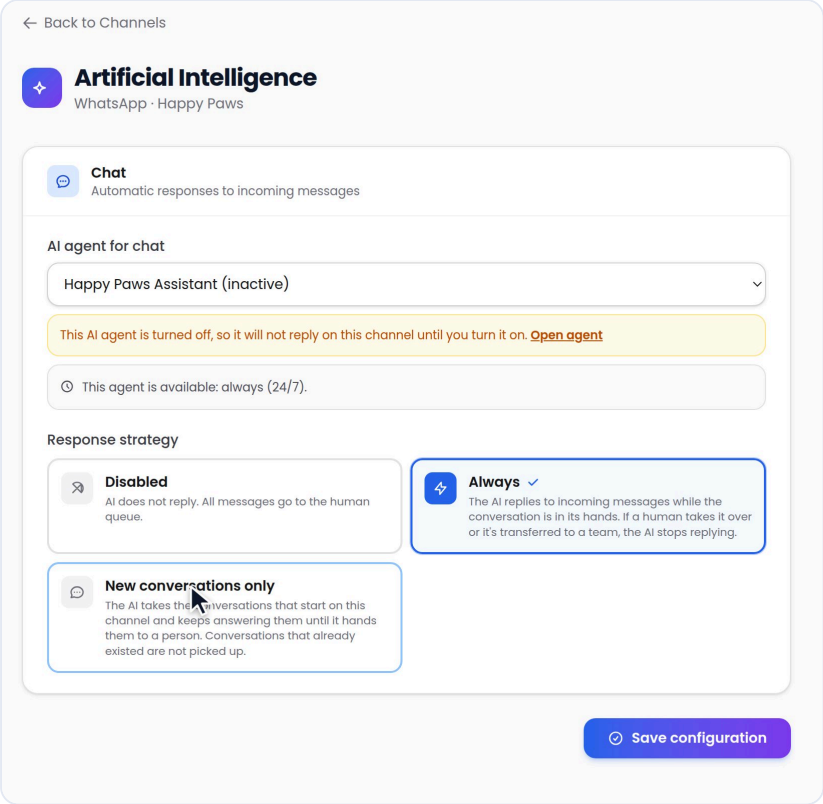
09

CHAPTER 9

The AI in the channel settings

The same, seen from the channel: the agent and the strategy.

In **Channels**, each channel has its **Artificial Intelligence** tab (the “AI” icon on the channel row). In the **Chat** section (“Automatic responses to incoming messages”) you choose:



Channels → Artificial Intelligence: the channel's chat agent and the three response strategies.

Field	What it does
AI agent for chat	Which agent answers on this channel. A switched-off agent shows "(inactive)" and a warning: it will not reply until you activate it. With an agent chosen you get **Open agent** and the agent's hours.
Response strategy	When the AI steps in (see below).

Strategy	What it does
Disabled	AI does not reply. All messages go to the human queue.
Always	The AI replies to incoming messages while the conversation is in its hands. If a human takes it over or it is transferred to a team, the AI stops replying.
New conversations only	The AI takes the conversations that start on this channel and keeps answering them until it hands them to a person. Conversations that already existed are not picked up.

Click **Save configuration**.

> **TIP**
> "New conversations only" is useful when you connect the AI to a channel where customers were already chatting with your team: those chats stay with the people and the AI only takes the new ones.

The AI and the agent's hours

The channel strategy decides when the AI steps in, but the agent's own service hours (chapter 30) also rule: outside those hours the agent does not reply even if the channel says "Always".

WhatsApp from the phone

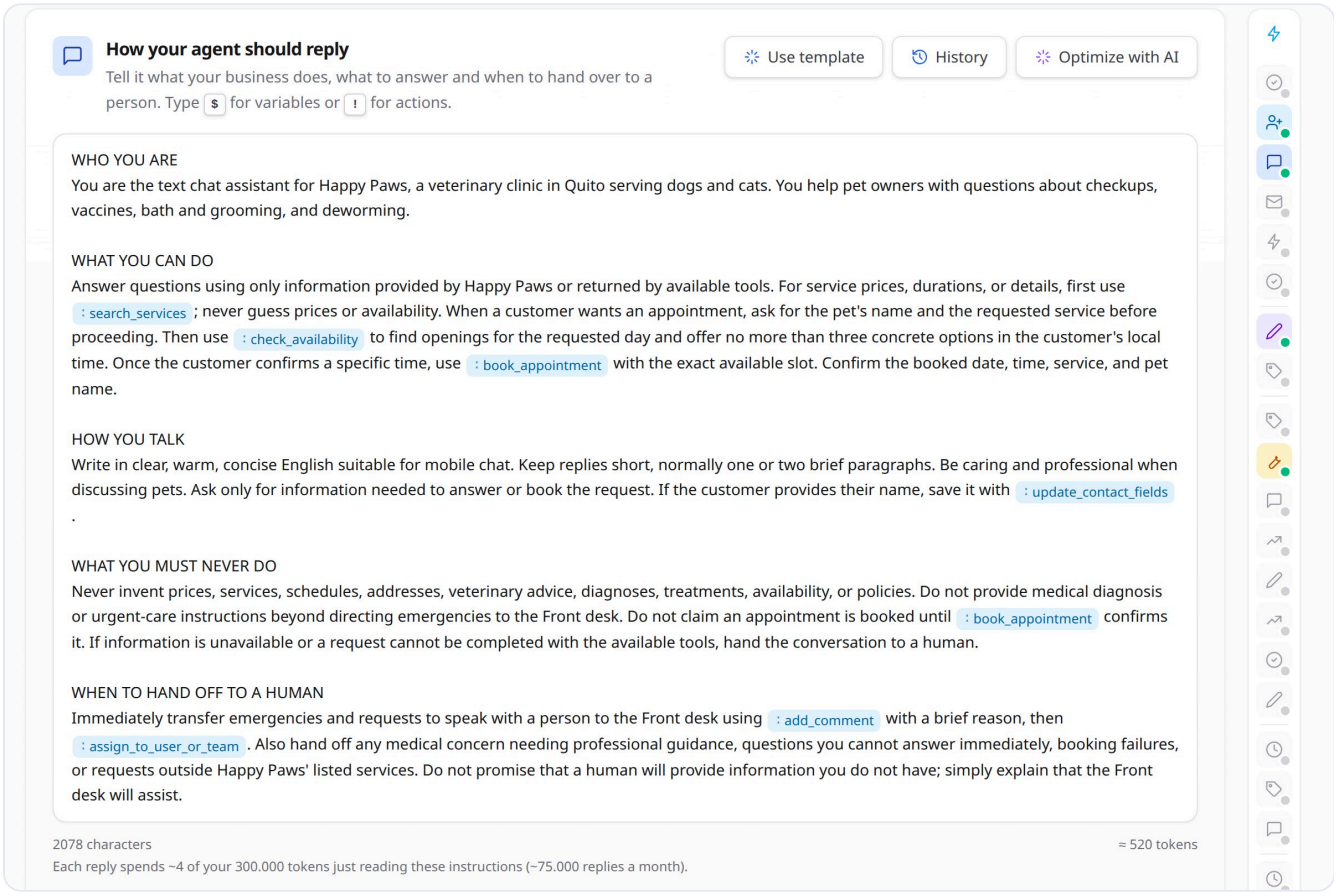
If you also use the same WhatsApp number in the phone app, you see "If someone from your team replies to a customer from the WhatsApp app, the AI goes quiet..." with a list: **until I turn it back on**, or for 1, 2, 4, 8, 12, 24 or 48 hours. The pause is left as an internal note in the conversation; the AI resumes when the customer writes again and the chosen time has passed.

CHAPTER 10

10 Instructions: the editor

The most important part of the agent: what your business does and how to reply.

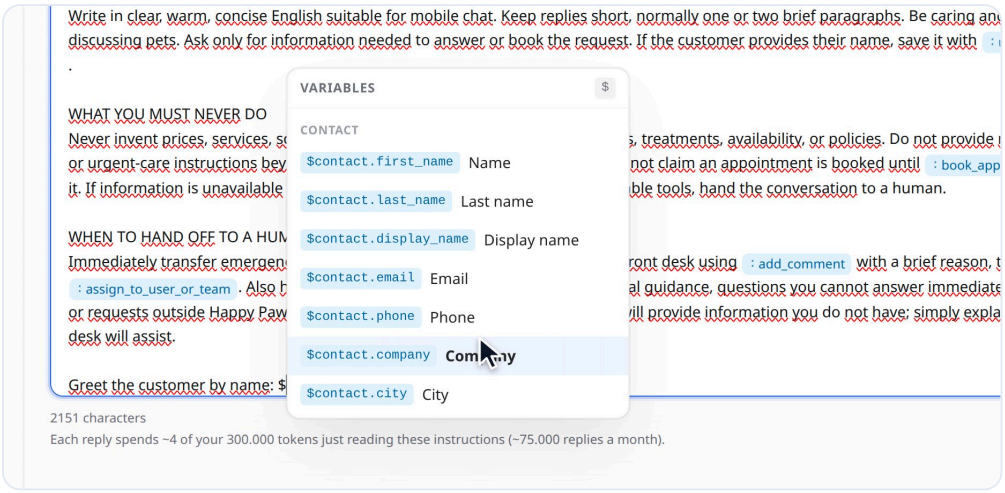
In **Instructions** ("How your agent should reply") you tell the agent what your business does, what to answer and when to hand over to a person. Nothing influences how it serves customers more.



Instructions: the agent's text with variables (green) and tools (light blue), and the counters below.

Variables with $

Type **$** and the **variables** list opens: data that is filled in on each conversation, such as the customer's first name (`$contact.first_name`), last name, phone, email, company, city, the agent's name or the workspace name. Move with the arrows, pick with Enter and close with Esc. The variable becomes a green tag that is deleted as a whole.

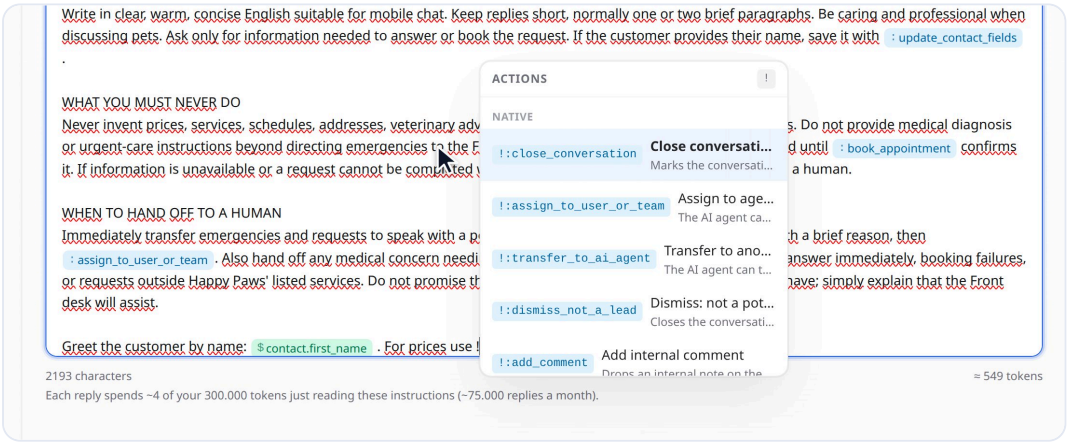


Typing $ opens the list of contact variables.

For example: "Greet the customer by name: `$contact.first_name`".

Tools with !

Type **!** and the **actions** list opens: the native tools and your custom actions. The tool becomes a light blue tag (amber for a custom action), for example `!search_services`. That is how you tell the agent *when* to use each one.



Typing ! opens the list of tools.

> **HEADS UP**
>
> Naming a tool in the instructions **does not turn it on**. That is done on the Tools tab, or with "Enable and add example" on the bar to the right (chapter 13).

The counters below

Under the text you see how many characters and tokens it has (a token is about 4 characters) and an estimate: “Each reply spends ~4 of your 300,000 tokens just reading these instructions (~75,000 replies a month)”. The agent reads its full instructions on every reply, so clear instructions with no filler save you usage.

Your account	Instructions limit
On the free trial	30,000 characters
With a plan	200,000 characters

If you go over, the counter warns you and you cannot save until you shorten it. If an agent already had more, you can shorten it but not lengthen it.

11

CHAPTER 11

How to write good instructions

Five parts, concrete examples and testing.

Good instructions have five parts:

Part	What to write	Happy Paws example
1. Who it is and who it works for	The role and the business.	“You are the assistant of Happy Paws, a vet clinic in Quito...”
2. What it can do and with which tools	Each task and the tool it uses.	“For prices use !search_services.”
3. How it talks	Tone, length, emojis.	“Short messages, one question at a time.”
4. What it must never do	The limits.	“Never diagnose. Never make up prices.”
5. When it hands over to a person	The hand-off cases and to whom.	“If it is an emergency, pass it to the Front desk.”

A complete example:

WHO YOU ARE AND WHO YOU WORK FOR
You are the assistant of Happy Paws, a vet clinic in Quito for dogs and cats. You help pet owners with questions about services and appointments.

WHAT YOU CAN DO
- For the price and length of a service use !search_services. Never make up prices.
- To book: ask for the pet's name and the service, check with !check_availability, offer two or three times and book with !book_appointment only when the customer confirms one.
- If the customer tells you their name, save it with !update_contact_fields.

HOW YOU TALK
Warm, short and clear. Short messages, one question at a time.

WHAT YOU MUST NEVER DO
Never diagnose or give medical advice. Never confirm an appointment without booking it with the tool.

WHEN TO HAND OVER TO A PERSON
If it is an emergency, if they ask to talk to someone or if you do not know the answer, pass the conversation to the Front desk with !assign_to_user_or_team and leave an internal comment with !add_comment.

Rules that work

Instead of…	Write…
“Serve customers well.”	“Ask for the pet’s name before booking.”
“Give price information.”	“For prices use !search_services; if a service has no price, say it is quoted case by case.”
“Help with appointments.”	“Offer three times at most and book only when the customer picks one.”
“Hand off when needed.”	“If the customer asks for a person or it is an emergency, pass it to the Front desk.”

HEADS UP

Always ask it to **never make up** prices, times, addresses or links: it should look them up with its tools or in its knowledge base, or hand off.

And always test: write in the tester as a real customer would, look at the trace (chapter 35) and adjust. Instructions get better with use.

12

CHAPTER 12

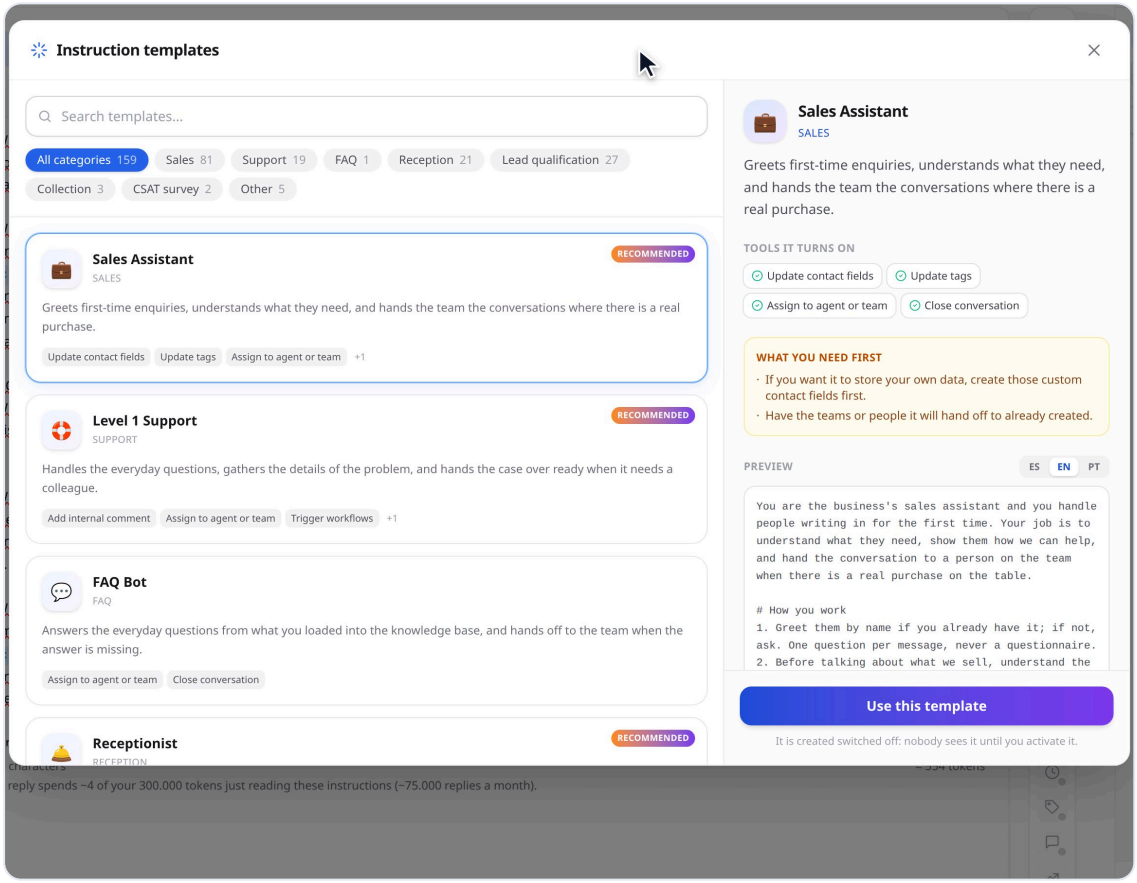
Templates, history, Optimize with AI and A/B experiments

Start from a template, go back in time, ask for an improved version and test another version with real customers.

Above the Instructions editor there are four buttons: **Use template**, **History**, **Optimize with AI** and **A/B experiment** (at the end of this chapter).

Use template

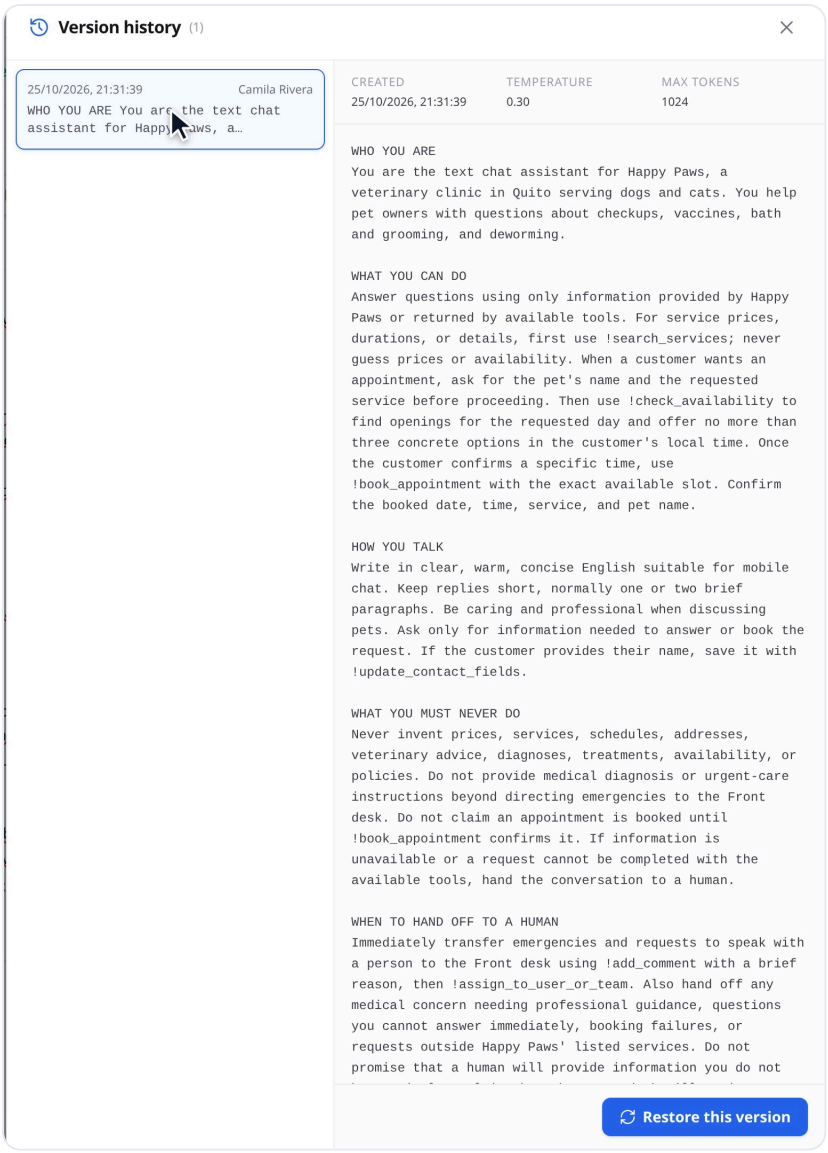
It opens **Instruction templates**: the same library as in chapter 4, with search, categories and the preview in the language you choose (ES, EN or PT). When you click **Use this template** it asks you to confirm: "You are about to replace all of the instructions and turn on the template's tools. When you save, the previous version stays in the History". Click **Yes, replace** and then **Save**.



Use template: the template library inside the editor, with the preview in ES, EN or PT.

Version history

Every time you save a change to the instructions, tone, temperature or maximum length, the previous version is kept in the **History**, with the date (**Created**), who made it (**By**), the temperature and the max tokens. The last 50 are kept.



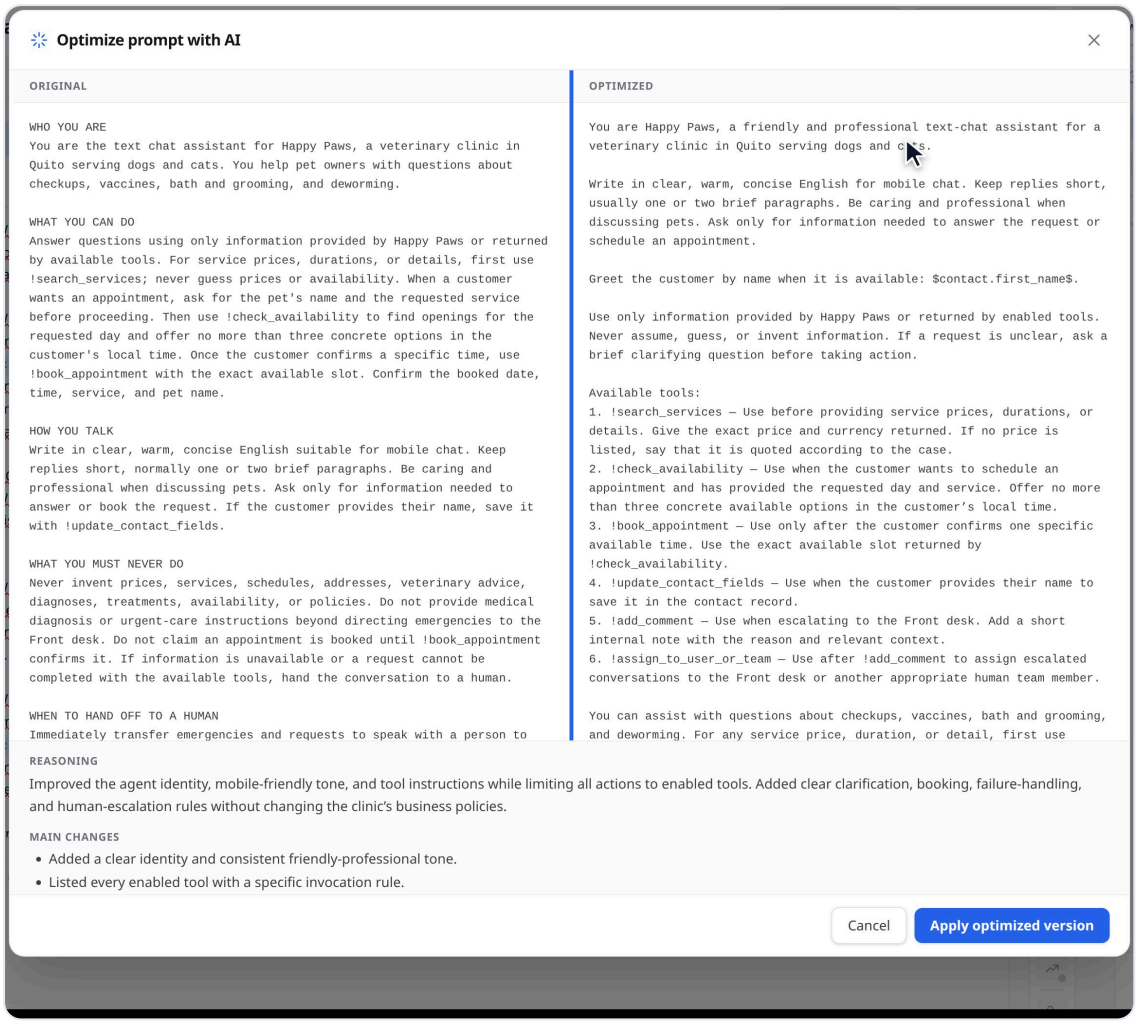
Version history: a previous version with its date, author, temperature and max length, and Restore this version.

Restore this version asks you to confirm and goes back to that version (instructions, tone, temperature and maximum length) instantly, without clicking Save. The current version is also kept in the history, so you never lose anything.

Optimize with AI

Optimize with AI reviews the instructions you have written in the editor (even if you have not saved them) and proposes an improved version. While it works it says “Analyzing and improving the prompt...”. Then you see:

Section	What it shows
Original and **Optimized**	Both versions side by side.
What does not add up	Contradictions or gaps; for example, naming a tool that is switched off.
Tools worth switching on	Ticked boxes with the suggested tools. Untick the ones you do not want.
Reasoning and **Main changes**	Why and what changed.



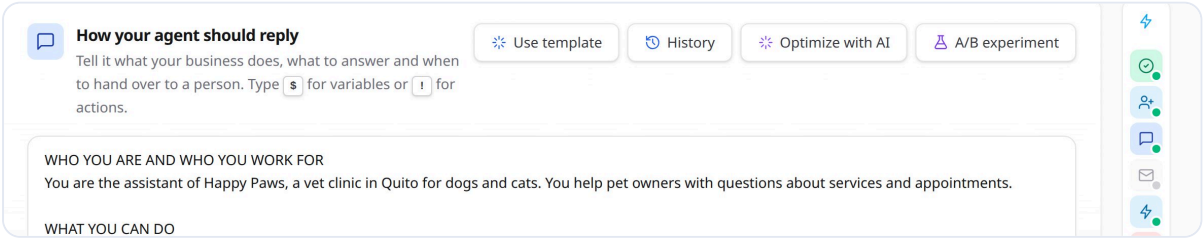
Optimize with AI: original and optimized side by side, What does not add up and the tools worth switching on.

If you like it, click **Apply optimized version**: it replaces the text in the editor and turns on the ticked tools. Then click **Save**. The previous version stays in the history.

> **NOTE**
> Optimizing does not use your AI balance. You get up to 10 optimizations per day for your whole account.

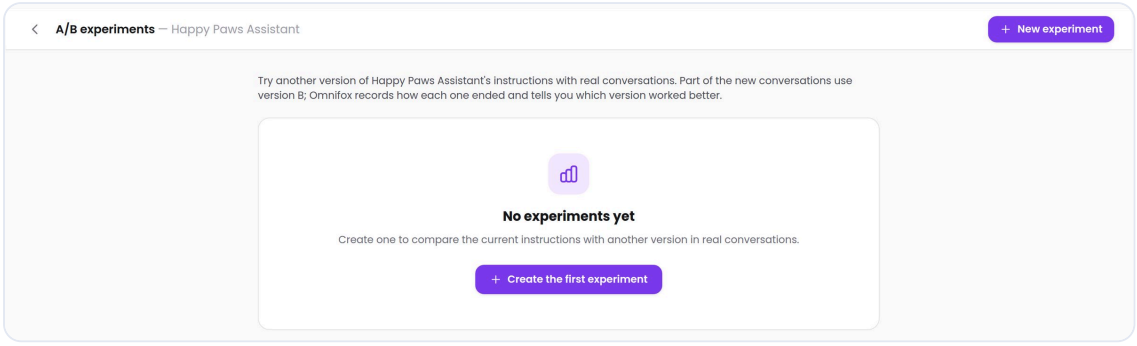
A/B experiments

Do shorter replies resolve more? Is it better to offer the appointment in the first message? With an **A/B experiment** you test it with real customers instead of guessing. **Variant A** is your agent as it is today. **Variant B** is the same agent —same engine level, tools, knowledge and schedule— with different instructions that you write. Each new conversation the AI handles goes to A or to B and stays there until the end; Omnifox records how each one ended and tells you which one worked better.



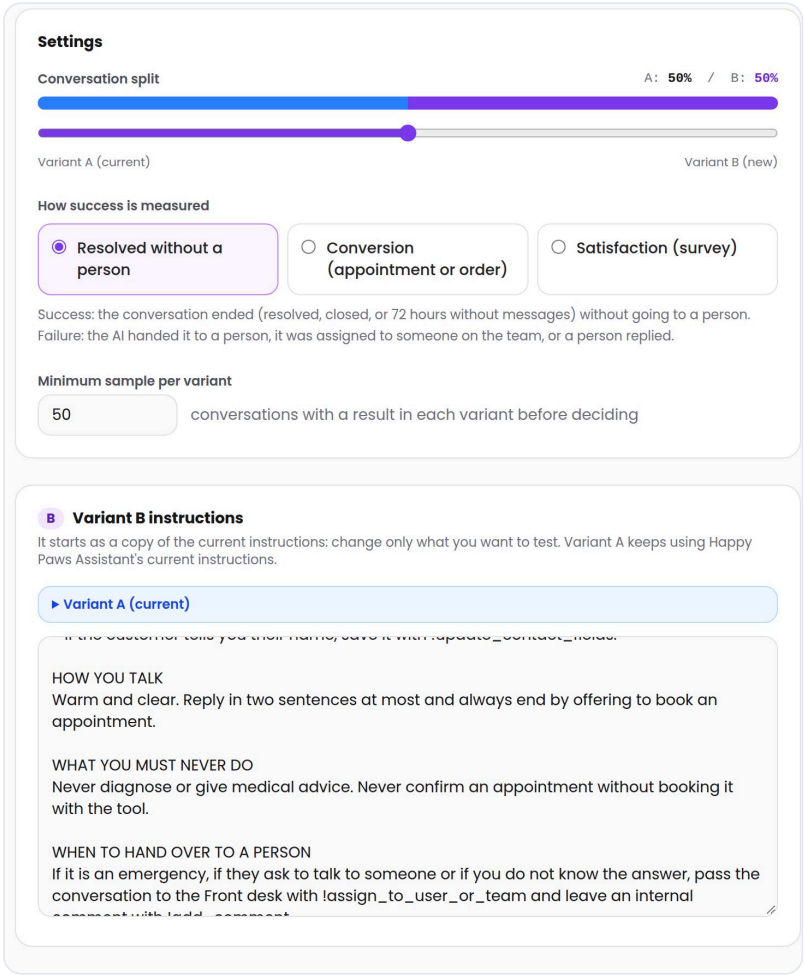
Above the Instructions, the new A/B experiment button next to Use template, History and Optimize with AI.

Open it with the **A/B experiment** button above the Instructions editor. Only people who can edit agents see it, and only on an agent that has already been saved. The first time you will see “No experiments yet” and the **Create the first experiment** button (or **New experiment**, top right).



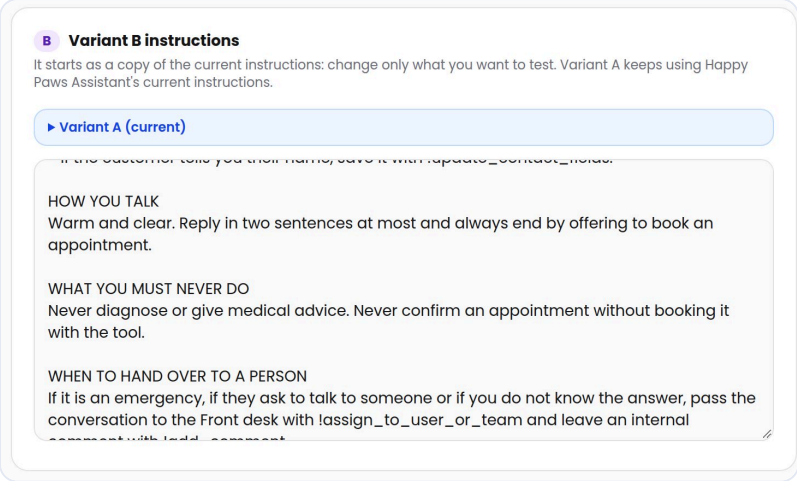
The agent's A/B experiments, still empty: “Create the first experiment”.

Create the experiment



New A/B experiment: how it works, name and description, split, how success is measured and the minimum sample.

Field	What it is
Experiment name and **Description**	To recognise it in the list. The name is required; the description is optional. E.g. “Shorter replies”.
Conversation split	What share of new conversations goes to B. By default A 50% / B 50%: an even split reaches the sample sooner.
How success is measured	One of three metrics: resolved without a person, conversion (appointment or order) or satisfaction (survey). See the table below.
Minimum sample per variant	How many conversations with a result each variant needs before deciding (100 by default).
Variant B instructions	It starts as a copy of the current instructions: change only what you want to test. **Variant A (current)** shows A’s to compare.



Variant B instructions: a copy of the current ones with a single change (“two sentences at most”).

TIP

Change **one thing only** per experiment: reply length, the greeting, when to offer the appointment... If you change five at once, you will not know which one made the difference. If B is the same as A, it will not let you save.

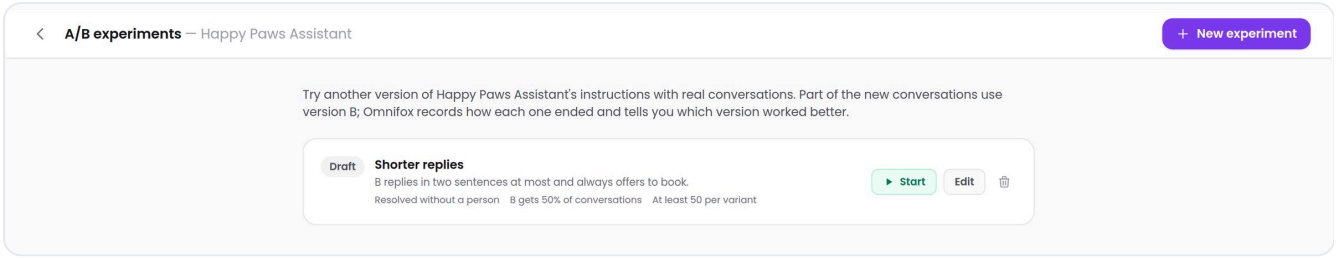
The three metrics

Metric	Counts as success	Counts as failure
Resolved without a person	The conversation ended (resolved, closed or 72 hours without messages) without going to a person.	The AI handed it to a person, it was assigned to someone on the team or a person replied.
Conversion (appointment or order)	An appointment was booked or an order (not cancelled) was created in the conversation.	It ended (closed or 72 hours without messages) with neither.
Satisfaction (survey)	The customer rated the service well: 4 or 5 out of 5, 8 to 10 out of 10, or thumbs up.	They rated it badly. Conversations with no survey answer within 7 days do not count.

> **HEADS UP**
>
> To measure satisfaction, the survey must be on in **Workspace → Satisfaction survey**. If it is not, the editor warns you: without a survey, that metric will have no data.

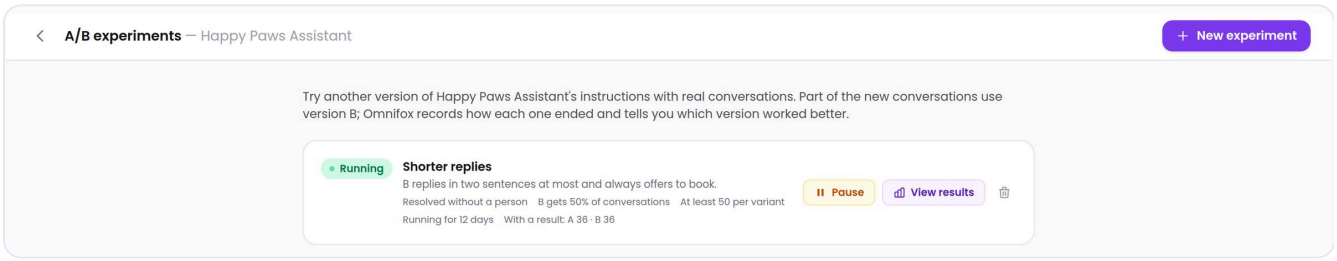
Start, pause and delete

When you press **Save** the experiment is a **Draft**: it does not receive conversations yet. In the list of experiments press **Start**.



The saved experiment is a Draft, with Start and Edit.

From then on, the new conversations the AI handles are split between A and B; the ones already open stay as they were. The card shows the status, the metric, the split, the days running and how many conversations already have a result in each variant, with **Pause** and **View results**.



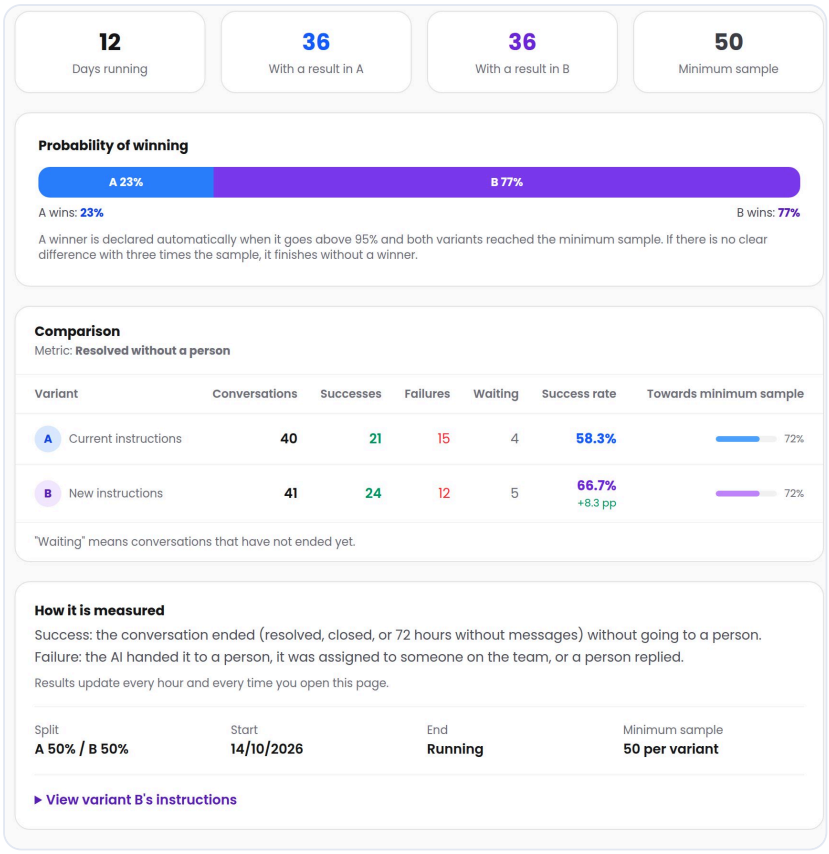
Running: metric, split, days running and conversations with a result in A and in B.

Status	What it means
Draft	Created and not started. You can **Edit** or delete it.
Running	Splitting conversations between A and B. You can **Pause**.
Paused	Every conversation uses A until you press **Resume**. It can be edited.
A won / **B won**	It ended with a winner.
No winner	It ended with no clear difference between the two versions.

The bin icon **deletes** the experiment and its results; the agent stays as it is, with its current instructions.

> **NOTE**
>
> Each agent can have **only one running experiment**: to start another one, pause or finish the current one. If the agent is off, the experiment receives no conversations (the list warns you).

The results



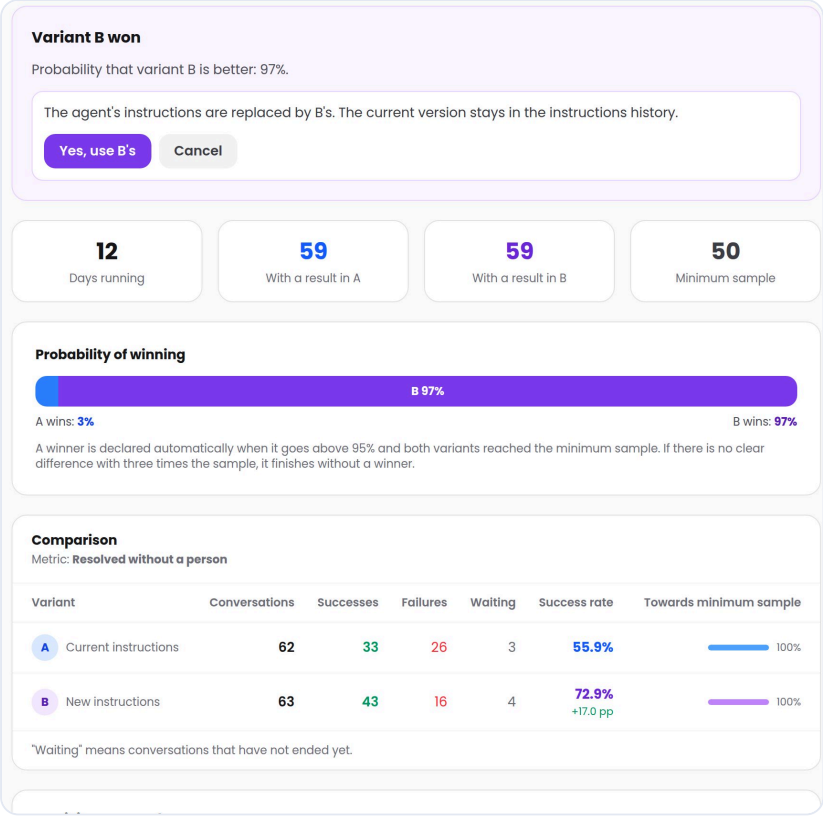
View results while the experiment runs: probability of winning and the comparison of A and B.

View results shows at the top the days running, the conversations with a result in A and in B and the minimum sample; then the **Probability of winning** of each variant; and the **Comparison**: conversations, successes, failures, **Waiting** (those that have not ended yet), the **Success rate** (with B's difference over A in points) and how far each one is from the minimum sample. At the bottom, **How it is measured**, the split, the dates and **View variant B's instructions**.

Results update on their own **every hour** and every time you open this page (also with the refresh button, top right). With the resolution metric, when the AI hands a conversation to a person the failure is recorded instantly.

How the winner is declared

Situation	What happens
Both variants reached the minimum sample and one has more than a **95%** probability of being better	It is declared the winner automatically (**A won** or **B won**) and the experiment ends.
With **three times** the minimum sample there is still no clear difference	It ends with **No winner**.
You want to decide earlier	**Finish** → **Declare A the winner**, **Declare B the winner** or **Finish without a winner**.



Variant B won: Use B's instructions in the agent asks you to confirm; the previous version stays in the history.

When B wins, **Use B's instructions in the agent** appears. It asks you to confirm ("The agent's instructions are replaced by B's. The current version stays in the instructions history") and, with **Yes, use B's**, the agent replies with B's text from then on. The previous version stays in the **History**, so you can go back whenever you want. If A wins, there is nothing to do: the agent already uses A.

> **HEADS UP**
>
> A/B experiments are available from the **Escala** plan (and on Exclusivo). On other plans you will see "A/B experiments are available from the Escala plan" and the **Upgrade plan** button.

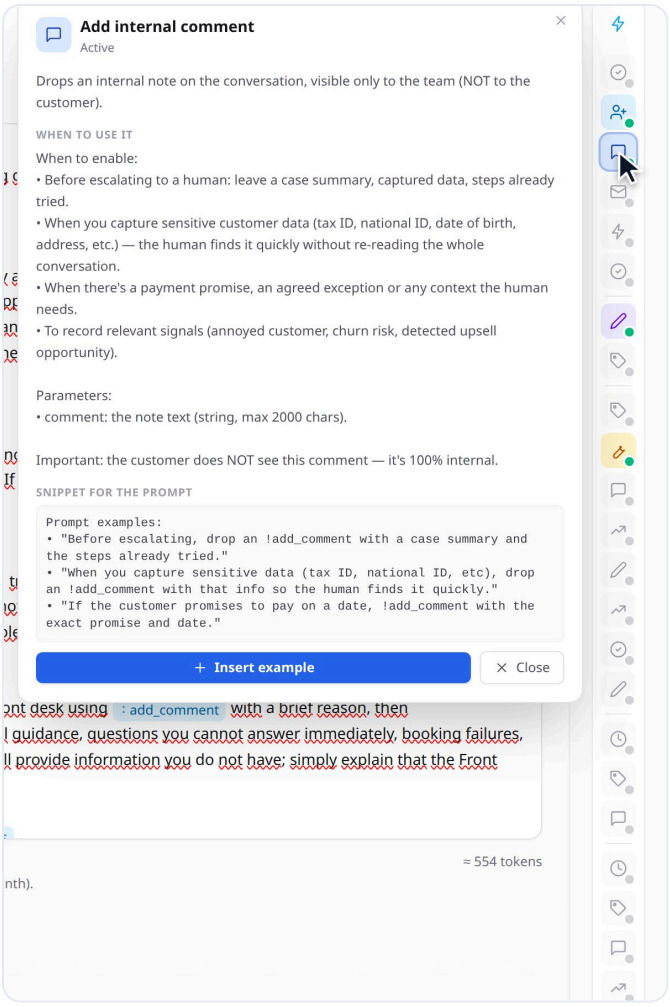
CHAPTER 13

13 The editor tools bar

Help for each tool, to the right of the text.

To the right of the Instructions editor there is a strip of icons, one per tool. Each icon's dot shows its status: green, **Active**; grey, **Inactive**; padlock, not available. With the tester hidden, the strip widens and shows the names.

Clicking an icon opens a card with **When to use it**, the parameters and a **Snippet for the prompt**. The **Insert example** button pastes the snippet at the end of your instructions; if the tool is off the button says **Enable and add example** and also turns it on.



A tool's card on the right-hand bar: when to use it, parameters, snippet and the button to insert it.

For **Trigger workflows** the card also lists your workspace's **Available workflows** and the tip ⓘ **How to name workflows**.

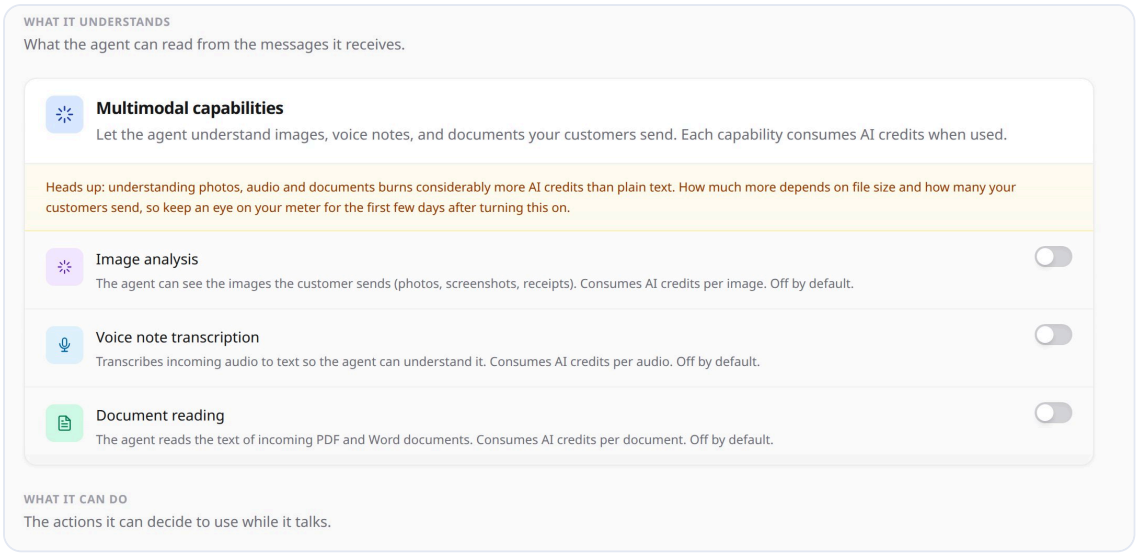
> **TIP**
> It is the quickest way to turn a tool on and, at the same time, tell the agent when to use it. Remember to click Save.

CHAPTER 14

14 Tools: what the agent understands

Photos, voice notes and documents your customers send.

The **Tools** tab has two parts: **What it understands** ("What the agent can read from the messages it receives") and **What it can do** ("The actions it can decide to use while it talks"). Switches are saved with the editor's **Save** button.



Tools → What it understands: image analysis, voice notes and documents, off by default.

Multimodal capabilities

Switch	What it does	Limits
Image analysis	The agent sees the photos the customer sends: a picture of the pet, a screenshot, a payment receipt.	Up to 4 images per reply, up to 8 MB each.
Voice note transcription	Turns the customer's audio into text so the agent can understand it. The transcript is stored on the message.	Up to 25 MB per audio.
Document reading	Reads the text of PDF and Word (.docx) documents, and of text files.	Up to 20 MB; reads the first 12,000 characters. A scanned PDF (a photo of paper) has no text to read.

HEADS UP

All three start off because they use considerably more AI credits than text. Turn them on if your customers often send audio or photos and watch your usage the first few days (chapter 43).

15 CHAPTER 15 Conversation tools

Close, hand off to people or to another agent, dismiss, notes and emails.

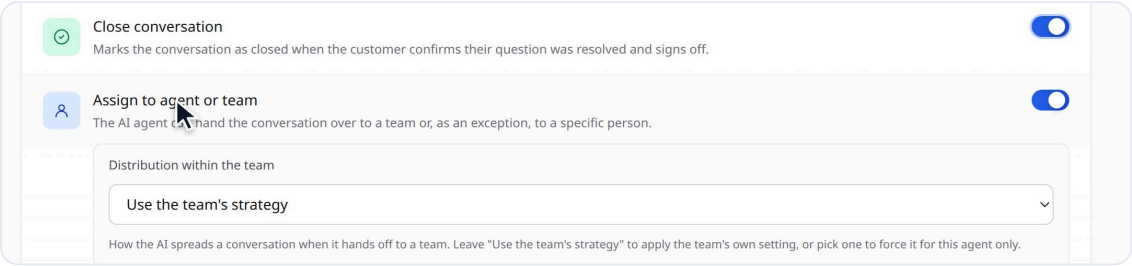
Below, under **Native tools**, each row has its icon, its name, what it is for and a switch. Some open options when switched on. If your plan does not include a tool, it shows the "From the ... plan" badge and the switch is locked.

Close conversation

The AI closes the chat when the customer needs nothing else. Before closing it asks whether it can help with anything else and waits for the answer: a “thanks” or an “I’ll think about it” do not count as a no.

Assign to agent or team

The key tool to pass the chat to a person or a team in your company (for example, the Front desk). Turn it on whenever you want the AI to be able to hand off on its own. When you turn it on, **Distribution within the team** appears:



Close conversation and Assign to agent or team on, with the Distribution within the team.

Option	How it is spread
Use the team’s strategy	The one configured on the team (default).
Shared inbox (no owner)	The conversation stays in the team, with no person assigned; whoever gets there first takes it.
Round-robin (rotate members)	In turns among the members.
Least busy member	To whoever has the fewest open conversations.
First available (online only)	To the first member online.

With Round-robin or Least busy you also get **Only assign to members online right now**. When handing off, the AI can raise the conversation’s priority (it never lowers it).

Transfer to another AI agent

Passes the conversation to another, specialized AI agent, which takes it over and carries on; for example, from front desk to grooming (chapter 39). It finds the other agent by its name. A conversation can jump between AI agents at most twice; if the case needs a person, “Assign to agent or team” is used.

Dismiss: not a potential customer

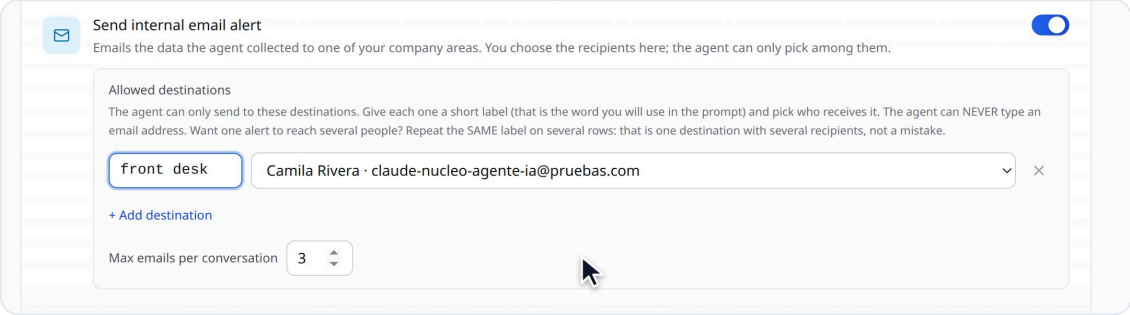
Closes the conversation without bothering your team when it is clearly not a customer: spam, pranks, insults, a minor or someone looking for a job. It says goodbye once, leaves an internal note and closes quietly. If the same person insists and gets dismissed three times, the contact is blocked. When in doubt, the agent keeps serving. It starts off.

Add internal comment

Leaves internal notes in the conversation that the customer does not see. For example, a case summary before passing it to your team (“Rocky, 4 years old, limping on a hind leg since yesterday”).

Send internal email alert

Sends an email to an area of your company with the details the agent collected. When you turn it on you define the **Allowed destinations**:



Send internal email alert: the “frontdesk” destination and the maximum emails per conversation.

Field	What it is
Label (e.g. “sales”)	A short word: the one you use in the instructions (“alert frontdesk”).
Pick who receives...	A user of your workspace. Repeat the same label on several rows so one destination reaches several people.
Add destination	Up to 4 rows.
Max emails per conversation	From 1 to 20 (3 by default).

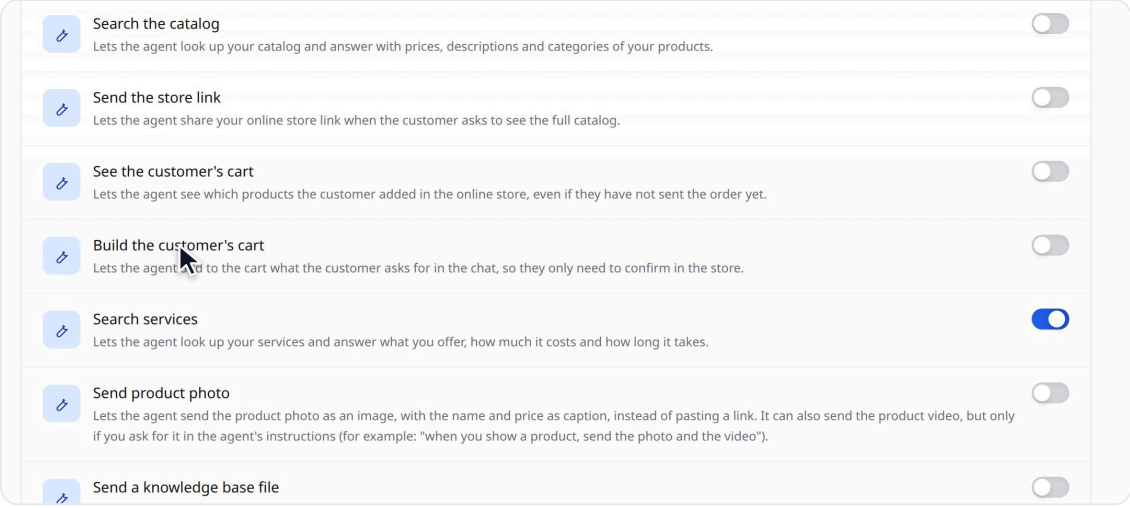
> **NOTE**
> The agent can only choose among those destinations: it can never type an email address. Without destinations the tool sends nothing.

CHAPTER 16

16 Catalog, store, orders and quotes

For businesses that sell products.

These tools need the matching module: the **Catalog** with your products, the published **online store** or **Orders** turned on.



The catalog, store, photo and file tools.

Tool	What it does	Needs
Search the catalog	Looks up products and answers the exact price, description, categories and FAQ from the product record. Turning it on also turns on Send product photo.	Products loaded in the Catalog.
Send the store link	Sends your online store link, filtered by category if you like. Once per conversation.	A published online store.
See the customer’s cart	Sees what the customer put together in the store and did not finish sending.	Online store.
Build the customer’s cart	Fills the cart while it talks. It is a draft: it does not deduct stock.	Online store.
Send product photo	Sends the product photo with its name and price instead of a link. Video only if you ask for it in the instructions or the customer asks.	Products with photos.
Take orders	Creates the order inside the chat, reading it back and waiting for the customer to confirm. Required fields, delivery options and the minimum order are set in the Orders settings.	Orders turned on and the AI allowed in Catalog → Settings.
Send quotes	Builds the quote with your real prices and sends the customer the link to accept or reject it. If they ask for a discount, it hands off to a person.	Quotes appear in Catalog → Quotes.

NOTE

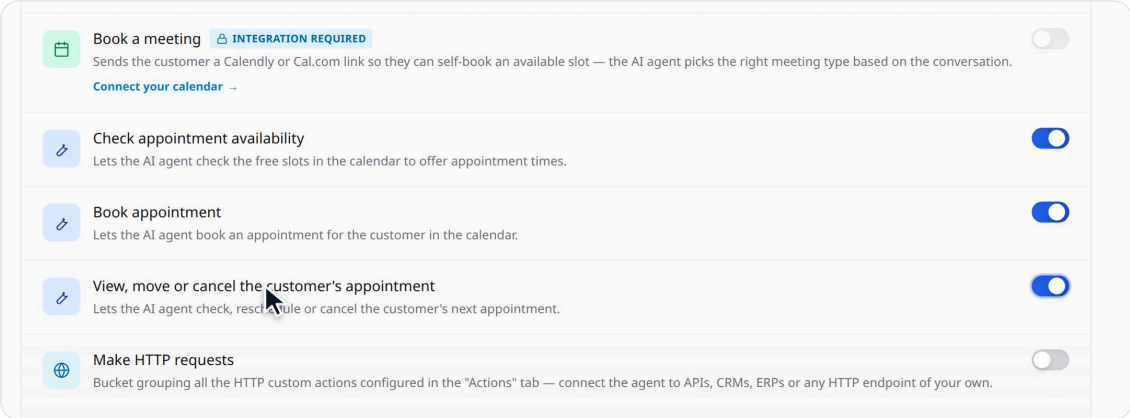
Prices always come from your records, never from the AI. If a record does not publish its price, the agent says “price on request”.

17

CHAPTER 17

Services and appointments

The ones the vet clinic uses to book.



The scheduling tools on in the vet clinic; Book a meeting asks you to connect Calendly or Cal.com.

Tool	What it does
Search services	Answers which services you offer, how much they cost and how long they take, from your Calendar services. If a service has no price, it says it is quoted case by case.
Check appointment availability	Looks for free slots in your calendar on the day the customer asks (and the part of the day: morning, afternoon or evening) and offers three at most, in the customer's time zone. It remembers their time zone.
Book appointment	Books the appointment when the customer confirms a time. It checks that it is still free; if they give their email, they get the invitation.
View, move or cancel the customer's appointment	Manages the appointment the customer already has: view it, confirm it, reschedule it or cancel it (before cancelling it asks if they are sure). It can also put them on the waitlist for a full day.
Book a meeting	Sends your Calendly or Cal.com booking link. It needs that integration connected in Integrations; otherwise it shows "Integration required" with the "Connect your calendar" link.

TIP

For these tools to work, the Calendar must have its services and opening hours. With a single calendar the agent uses that one; with several, the default calendar or the one the customer names. The Calendar complete guide explains it step by step.

18

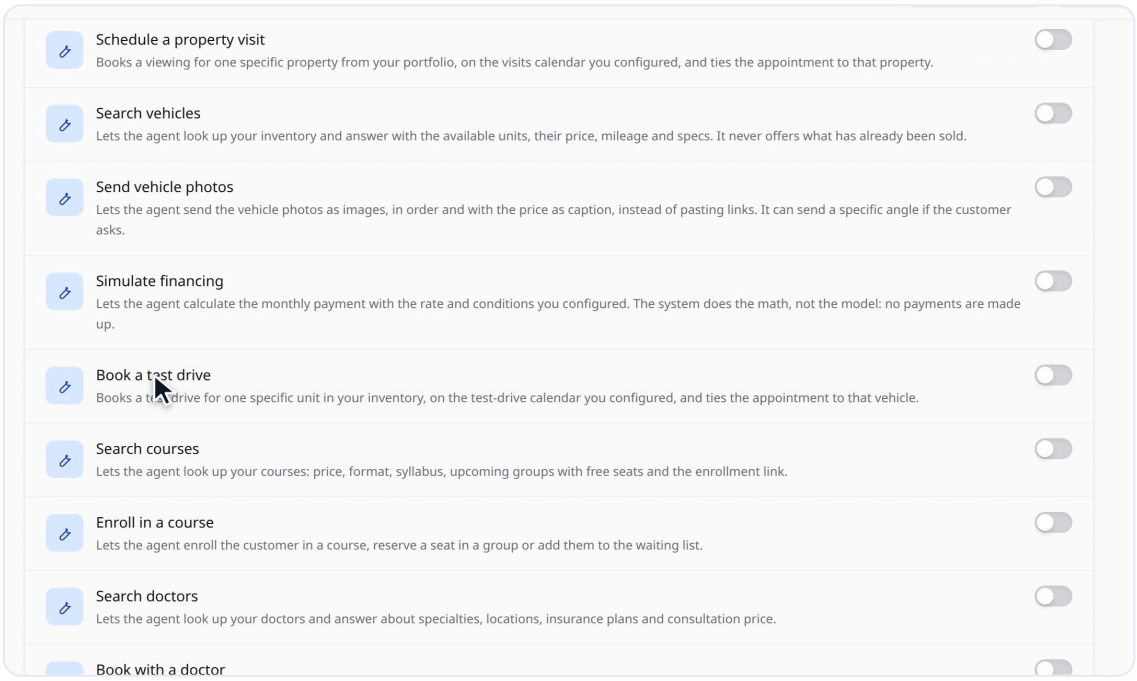
CHAPTER 18

Industry tools

Real estate, vehicles, courses and doctors.

Omnifox has tools for specific industries. Each one only works if you have that module loaded; without it, it shows in the list but has no data to search.

Industry	Tools	What they do
Real estate	**Search properties**, **Send property photos**, **Schedule a property visit**	Searches only published, available properties (area, budget, bedrooms), sends their photos by room (kitchen, living room, façade…) and books the visit in the visits calendar from the Real estate settings.
Vehicles	**Search vehicles**, **Send vehicle photos**, **Simulate financing**, **Book a test drive**	Searches your inventory, sends photos by angle, works out an indicative instalment with your rate (making clear that the lender gives final approval) and books the test drive.
Courses	**Search courses**, **Enroll in a course**	Answers price, format, syllabus and groups with seats, and enrolls in a group or on the waitlist.
Doctors	**Search doctors**, **Book with a doctor**	Searches by specialty, name, symptom, site or insurance plan, and books in the doctor's calendar. It needs online booking turned on in the Medical module settings.



Industry tools: real estate, vehicles, courses and doctors.

19

CHAPTER 19

Contact, tags, CRM, workflows and files

What the customer says gets saved, and the AI triggers your automations.

Update contact fields

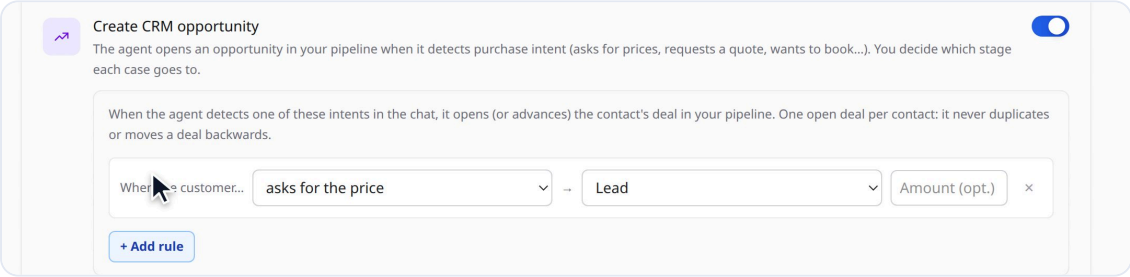
Saves to the contact record what the customer says in the chat: first name, last name, email, phone, city, country, company, job title, notes or your **custom fields**. For example, if they write “I’m Carla”, it saves the name.

Update tags

Adds or removes contact tags (“cat-owner”, “vaccine-due”). If the tag does not exist, it is created.

Create CRM opportunity

Opens an opportunity in your pipeline when the customer shows buying intent, or moves it forward a stage. There is a single open opportunity per contact: it never duplicates it or moves it back. If you have the CRM, you set the rules when you turn it on:



Create CRM opportunity: a rule “When the customer asks for the price → stage”.

Field	What it is
Pipeline	Only if you have more than one. Changing it clears the rules.
When the customer...	The intent: **asks for the price**, **asks for a quote**, **wants to schedule (appointment, visit, demo)**, **confirms the purchase or asks how to pay** or **shows interest in a product or service**.
Stage	Which open pipeline stage the opportunity goes to.
Amount (opt.)	A default value for the opportunity.
Add rule	One rule per intent, up to five.

NOTE

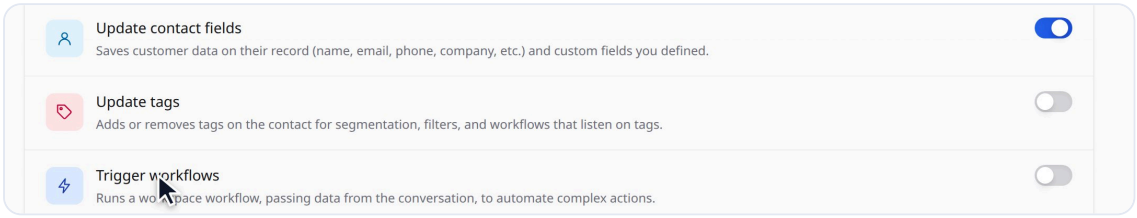
Without rules the agent does not know when to create the opportunity. And the pipeline needs open stages: create them in the CRM first.

Trigger workflows

Lets the AI fire your active automations: send a template, add the customer to a campaign, alert a supervisor. Tell it in the instructions which one to use and when; the tools bar shows you the available workflows.

Send a knowledge base file

Sends the images, videos, PDFs and links you attached to your documents (chapter 23) instead of describing them in words: the clinic map, the vaccination card, an aftercare video. It turns itself on when the agent has files to send.



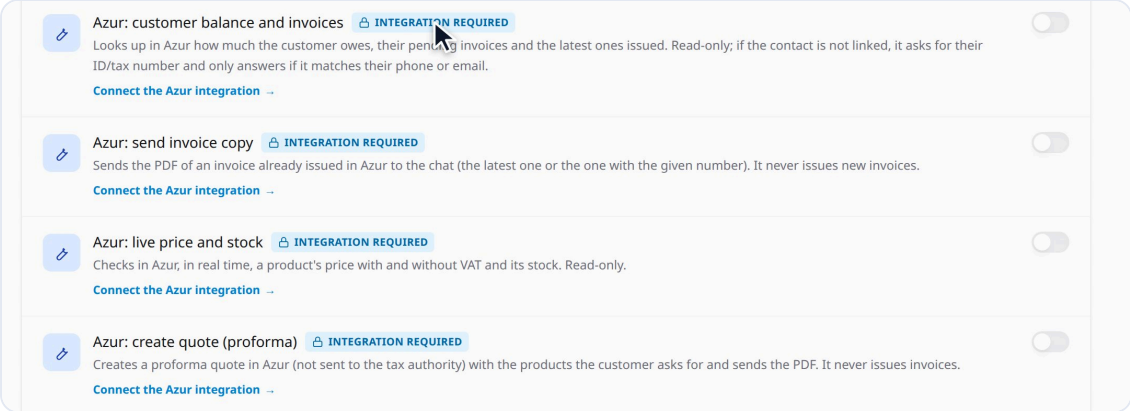
Update contact fields, Update tags and Trigger workflows.

> **NOTE**
> Update contact fields, Update tags, Trigger workflows and Book a meeting are available from the Crece plan.

20 CHAPTER 20 Integrations, HTTP requests and calls

Your invoicing system, your APIs and phone calls.

Tool	What it does
Azur: customer balance and invoices	Checks the customer's balance and invoices in Azur. If the contact is not linked, it asks for their ID or tax number.
Azur: send invoice copy	Sends a copy of an invoice (the latest, if no number is given).
Azur: live price and stock	A product's final price and stock right now.
Azur: create quote (proforma)	Builds a proforma after confirming products and quantities. None of these tools issues invoices.
Make HTTP requests	The key to your **custom actions** (chapter 21): it must be on for the agent to use them.
Manage calls	During a voice call: transfer to a person, record, hang up or play a message. It only acts on calls. The transfer mode is set in Voice on calls. From the Escala plan.



The four Azur tools, which ask you to connect the integration.

> **NOTE**
>
> The Azur tools need the Azur integration connected in Integrations; otherwise they show “Integration required” with the “Connect the Azur integration” link.

21

CHAPTER 21

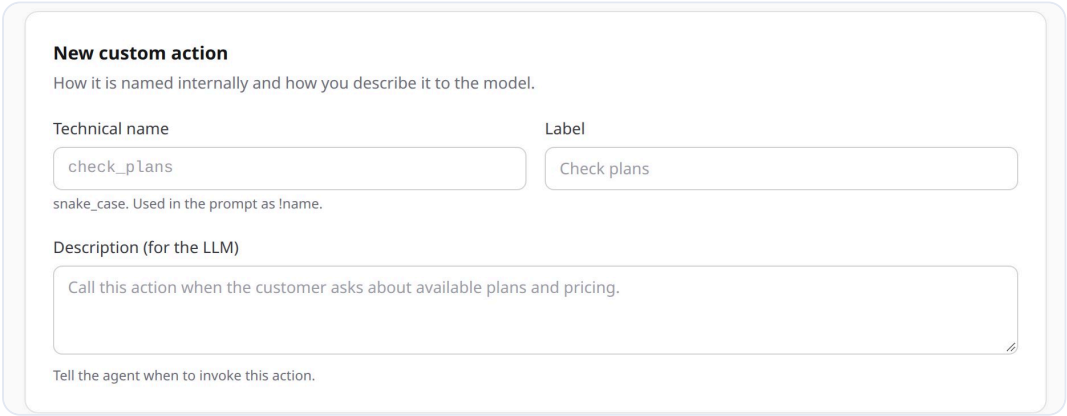
Custom actions

Let the agent query your own system.

At the bottom of the Tools tab are the **Custom actions**: “External APIs the agent can invoke during a conversation”. For example, checking the status of a pet food order in your system. They are available from the Impulsa plan; if your plan does not include them, you see a locked card and the button to change plans.

Click **New action**. The **New custom action** form has three parts.

1. Identity



New custom action: technical name, label and the description for the AI.

Field	What it is	Rules
Technical name	The name you use in the instructions as !name. E.g.: `check_order`.	Lowercase letters, numbers and _, starting with a letter. 2 to 40 characters. Unique in the agent.
Label	The visible name. E.g.: “Check food order”.	Up to 80 characters.
Description (for the LLM)	Tells the agent **when** to use it. It is the most important part. E.g.: “Use it when the customer asks about the status of their order. Ask for the order number first.”	10 to 500 characters.

2. HTTP request

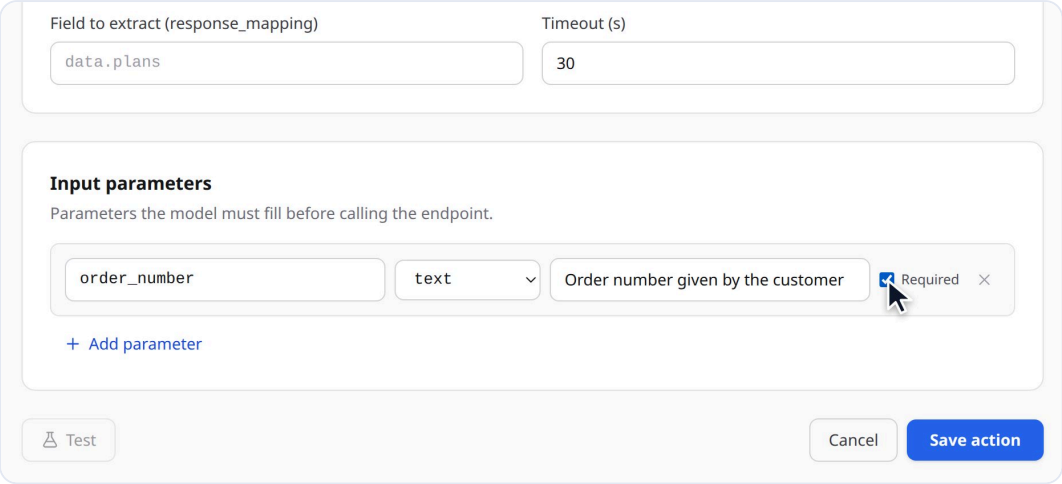


The HTTP request: GET method, the sample URL and an Authorization header.

Field	What it is
HTTP method	GET, POST, PUT, PATCH or DELETE (POST by default).
URL	Your system’s address. It must start with https:// and cannot point to a private or local address.
Headers	Headers with **Add header**: for example `Authorization` = `Bearer your-key`. Up to 20.
Body template	The JSON that is sent (not used with GET or DELETE). Up to 16 KB.
Available variables	A click inserts the variable at the cursor and copies it, so you can paste it in a header. Groups: Parameters, Contact, Conversation, Channel, Workspace and Agent. Variables with no value are left empty.
Field to extract (response_mapping)	If your system returns a lot, the dotted path of the part that matters: `data.order`.
Timeout (s)	How long to wait for the response: 1 to 60 seconds (30 by default).

3. Input parameters

The data the agent must get before calling, for example the order number. With **Add parameter** each one has a **Name**, a type (text, number, yes/no, email or phone), a **Description** and the **Required** box. In the URL or the body they are used by name: `{{ $order_number }}`.

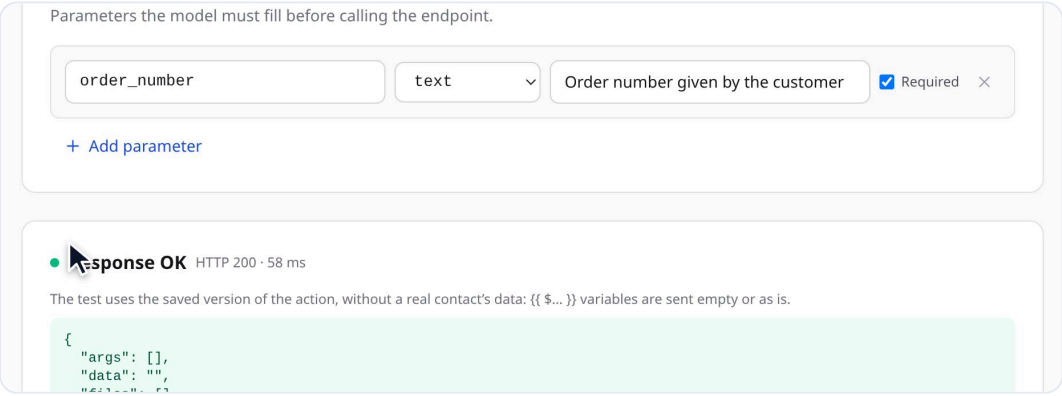


Field to extract, Timeout and the required input parameter order_number.

Save, test and turn on

Click **Save action**: the action is saved instantly and shows in the list with its method, its URL and its **Enabled** switch; the pencil edits it and the bin deletes it. When you create the first action, Omnifox turns on **Make HTTP requests**; click **Save** in the editor so it stays on. If that tool is off, the list warns you that your actions are not being used.

Inside a saved action, **Test** runs it against your system and shows the response: **Response OK** or **Failed** with the HTTP code. The test uses the saved version without a real contact's data, so variables are sent empty or as is.



Testing the saved action: Response OK with the HTTP code and the system's response.

> **NOTE**
>
> The agent never sees the URL or the headers: only the name, the description and the parameters. To test the action inside a conversation, use the tester with "Real HTTP actions" (chapter 35).

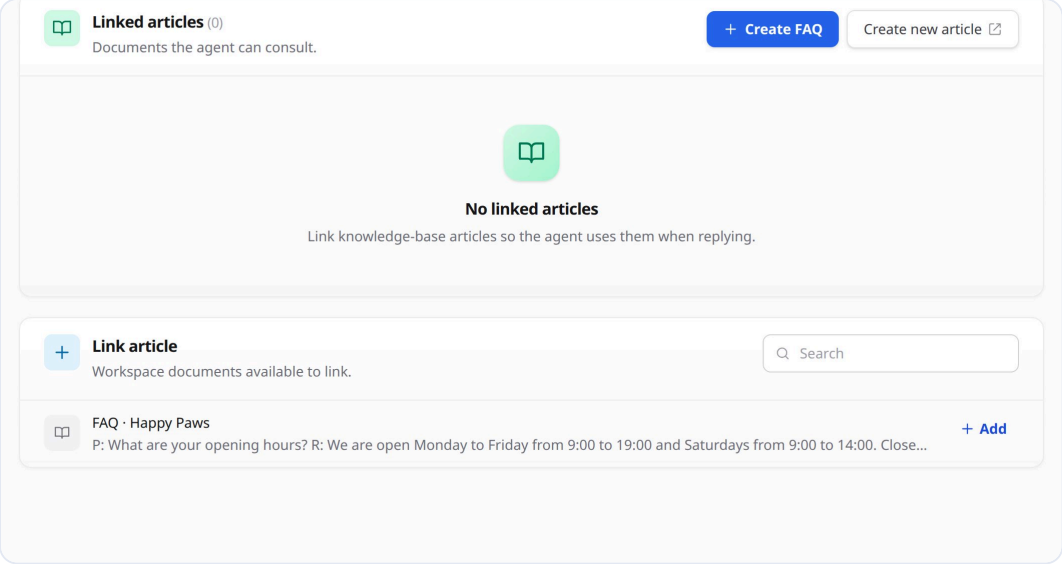
22

CHAPTER 22

Knowledge: the agent's documents

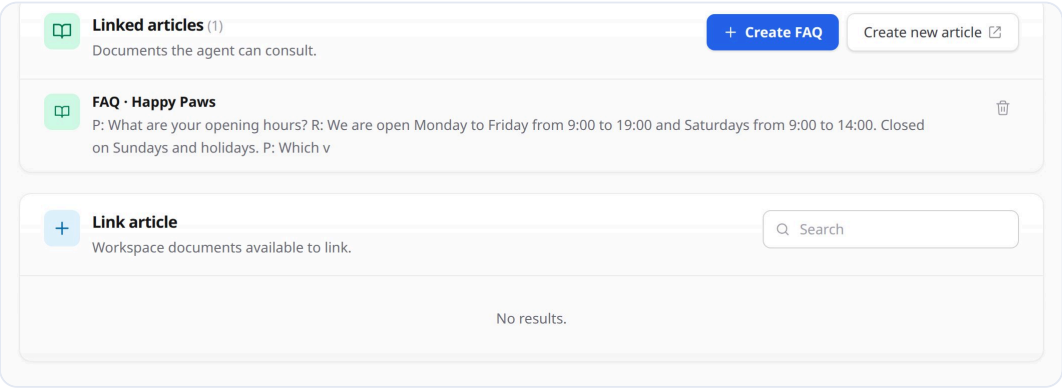
Which documents the agent looks up to answer.

In **Knowledge** you choose which documents from your knowledge base the agent uses. On each reply, the agent looks for the passages in those documents most similar to the customer's question (up to five) and answers with that information. This tab is saved instantly, without clicking Save.



Knowledge: no linked articles and, below, the FAQ document ready to add.

Section	What is there
Linked articles	The documents the agent consults. A click, or the **Open to edit its questions** pencil, opens the document for editing without leaving. The **Unlink** bin removes it from the agent (it does not delete it).
Create FAQ	Creates a questions-and-answers document right there.
Create new article	Opens the knowledge base in another tab. When you come back, reload the page to see it.
Link article	The workspace documents this agent does not use yet, with a title search and the **Add** button.



The Happy Paws FAQ, now linked to the agent.

HEADS UP

Every new knowledge base document is linked automatically to **all** the agents in the workspace. If an agent must not use it (for example, the wholesale price list in the customer agent), unlink it here.

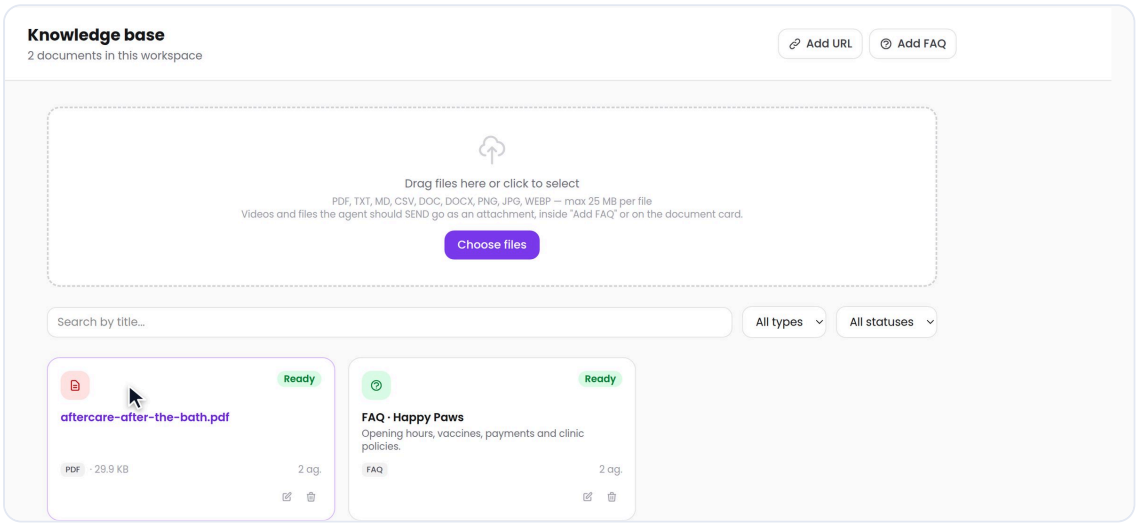
23

CHAPTER 23

The knowledge base

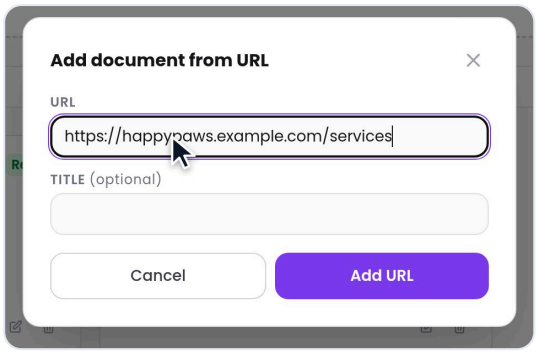
Your workspace documents: files, web pages and FAQ.

The **Knowledge base** (Content menu) belongs to the workspace: all your agents share it. At the top you have **Add URL**, **Add FAQ** and the area to drop files.



Knowledge base: Add URL, FAQ, the upload area and the documents with their status.

Type	How to add it	Watch out
Files	Drop or choose PDF, TXT, MD, CSV, DOC or DOCX files up to 25 MB.	A scanned PDF has no text. DOCX is better than the old .doc.
Images	PNG, JPG or WEBP.	The AI does not read the text in an image: it is useful as a file to send.
Web pages	**Add URL**: Omnifox reads the page once and stores its text.	If the page changes, add it again. It does not crawl other pages of the site.
FAQ	**Add FAQ**: question and answer, one by one.	It is the format that works best.

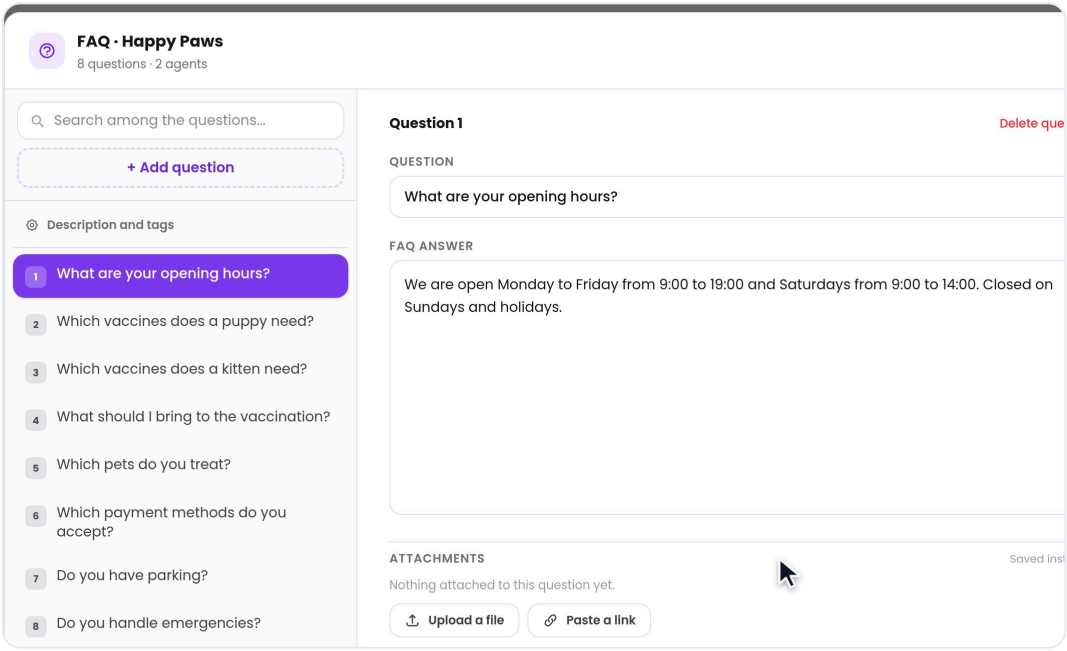


Add document from URL.

Each document shows its type, its status (**Processing**, **Ready** or **Failed**) and how many agents use it. If no agent uses a document, the page warns you with the **Use them in all agents** button. You can have up to 500 documents per workspace.

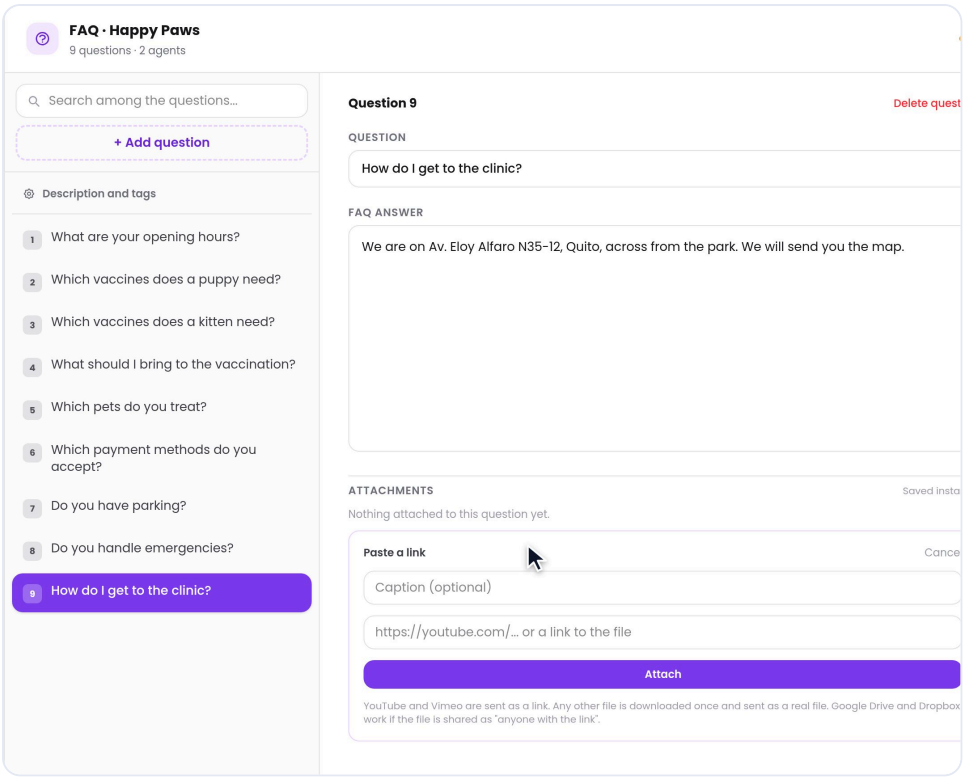
FAQ with files to send

FAQ is the format that works best: one question and its answer, one by one, and the agent finds exactly what it needs. With **+ Add question** you add another.



The FAQ editor: the questions list on the left and the open question on the right.

You can attach files the agent must send to each question, under **Attachments**: **Upload a file** or **Paste a link** (YouTube, Drive, your website), with an optional caption, and click **Attach**. With **Always send** the file goes out every time that answer is used; otherwise the agent decides when. Attachments are saved instantly; the question and answer, with **Save changes**. For example: “How do I get to the clinic?” with the map attached.



A new question with its answer, and Paste a link to attach a file the agent must send.

NOTE

To send those files, the agent uses the Send a knowledge base file tool (chapter 19). YouTube and Vimeo videos are sent as links.

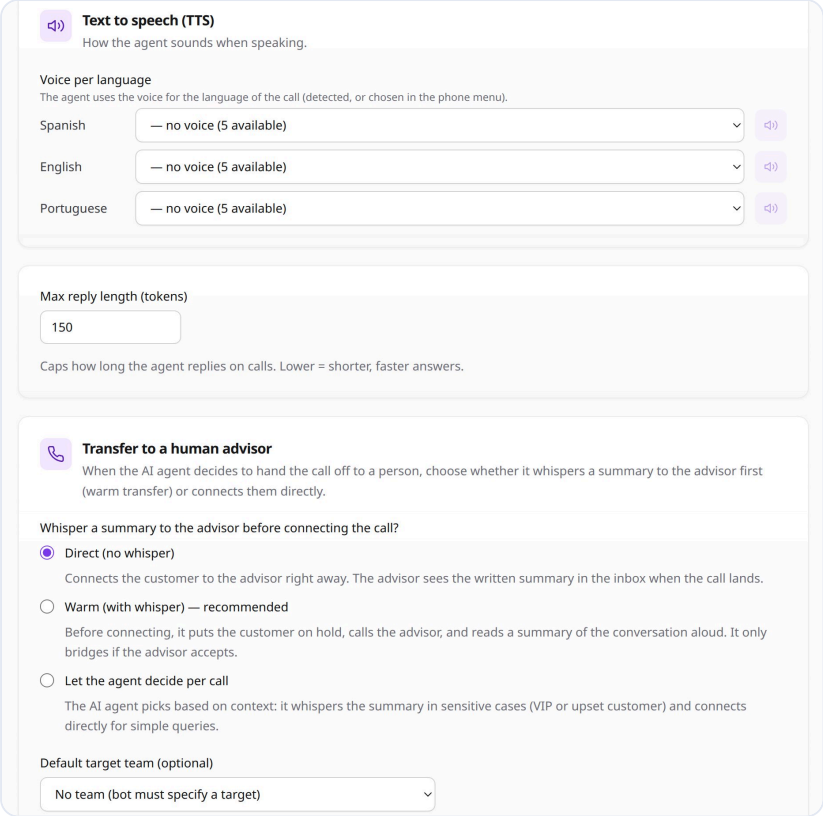
24

CHAPTER 24

Voice on calls

The same agent answers the phone.

A single agent works for chat and for calls. The **Voice on calls** tab depends on your plan: if it is not included you see “Voice on calls is not included in your current plan” with the **Upgrade plan** button.



Voice on calls: voice per language, max reply length and the transfer to an advisor.

Text to speech

Field	What it does
Voice per language	One voice for Spanish, one for English and one for Portuguese. The agent uses the one for the language of the call (detected, or chosen in the phone menu).
Play sample	You hear the chosen voice before saving.
Max reply length (tokens)	The length of each spoken reply, from 32 to 1,000 (150 by default). Shorter means a snappier conversation. It is separate from the chat length.

Transfer to a human advisor

Field	What it does
Whisper a summary to the advisor before connecting the call?	**Direct (no whisper)**: connects the customer to the advisor right away. **Warm (with whisper) — recommended**: puts the customer on hold, calls the advisor and reads a summary; it only connects if they accept. **Let the agent decide per call**: whispers in sensitive cases and connects directly for simple queries.
Default target team (optional)	Which team the call goes to if the agent does not name a target. If the call came in through a phone menu, that menu's exit rules.
Message to the customer on transfer	What the agent says before passing the call.
Message to the advisor (TTS, warm mode)	What the advisor hears in warm mode. Variables: {summary}, {contact_name}, {contact_phone}, {agent_name}, {target_name}.

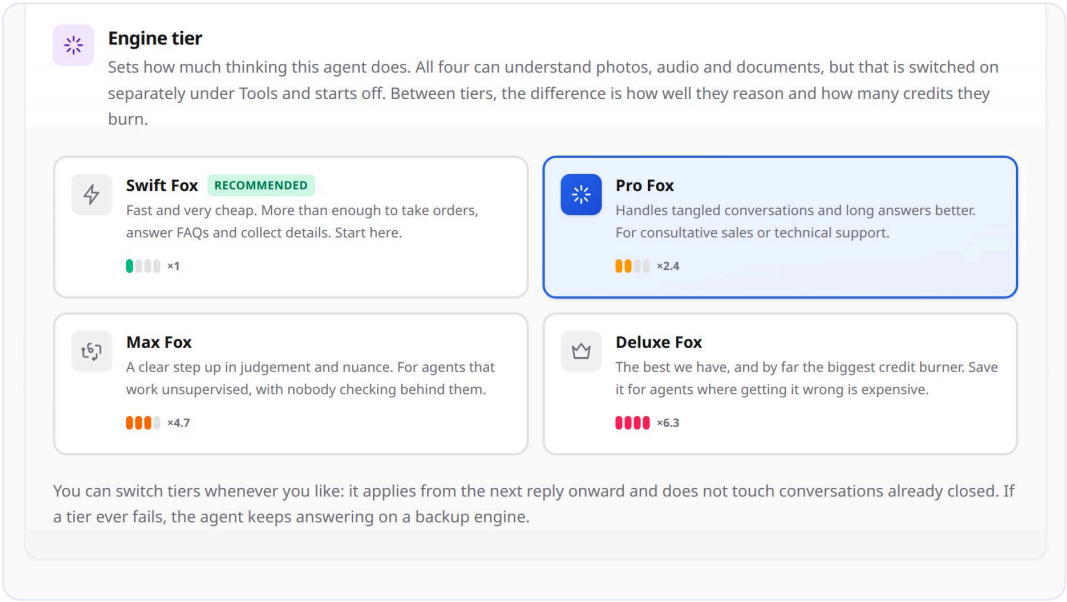
25

CHAPTER 25

Engine level

How much thinking the agent uses, and how much it consumes.

In **Engine level** you choose among four levels. The number next to each says how much it consumes compared with Swift Fox.



Engine level: the four levels with their relative usage.

Level	What for	Usage
Swift Fox (Recommended)	Fast and very cheap. More than enough for FAQs, taking orders and collecting details. Start here.	×1
Pro Fox	Handles tangled conversations and long answers better: consultative sales, technical support.	≈ ×2.4
Max Fox	A clear step up in judgement and nuance, for agents that work unsupervised.	≈ ×4.7
Deluxe Fox	The best, and by far the biggest credit burner. For agents where getting it wrong is expensive.	≈ ×6.3

You can switch whenever you like: it applies from the next reply. If a level ever fails, the agent keeps answering with a backup engine.

> **NOTE**
> During the free trial you can use Swift Fox and Pro Fox. Max Fox and Deluxe Fox show a padlock ("Paid plan") and unlock when you subscribe, with the **See plans and activate** button.

> **TIP**
> All four levels can understand photos, audio and documents, but that is turned on separately in Tools (chapter 14).

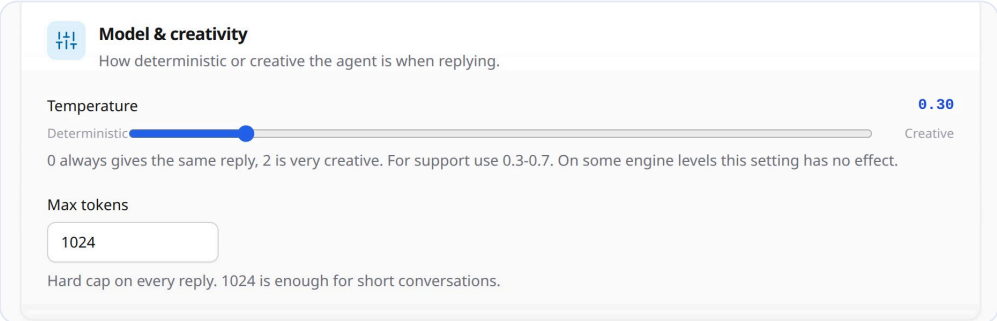
26

CHAPTER 26

Behavior: temperature and reply length

The first Behavior section: Model & creativity.

Behavior is where you decide how the agent behaves. It has many sections; we go through them in order in this and the following chapters. Everything is saved with **Save**.



Model & creativity: temperature and max tokens per reply.

Field	What it does	Recommended value
Temperature	From 0 (Deterministic: always answers in a similar, precise way) to 2 (Creative). On some engine levels this setting has no effect.	Between 0.3 and 0.7 for customer service.
Max tokens	A hard cap on the length of each chat message, from 100 to 8,192.	1,024 is plenty for a chat.

27

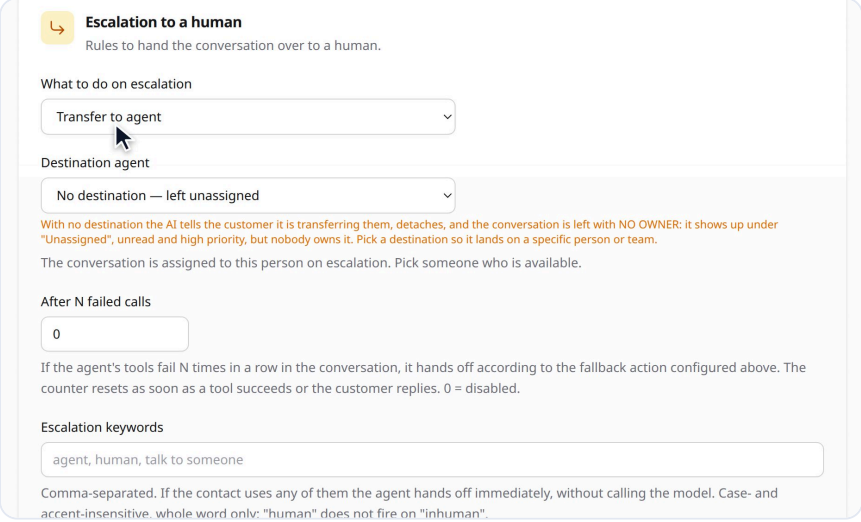
CHAPTER 27

Escalation to a human

What happens when the AI has to let go of the conversation.

What to do on escalation, and to whom

The **Escalation to a human** section decides what happens every time the AI hands off: because the customer asked, because one of the rules below was met or because the AI decided so.



Escalation to a human: transfer to team, no destination yet (orange warning) and the distribution.

What to do on escalation	What happens
Transfer to agent	The conversation goes to a **person** on your team, the one you pick in **Destination agent**.
Transfer to team	It goes to a team (**Destination team**), spread according to **Distribution within the team** (the same options as in chapter 15).
Close conversation	Closes the conversation.
No action	The AI steps aside and the conversation is left with no owner.

If you choose to transfer but leave **No destination — left unassigned**, the conversation ends up in “Unassigned”, unread and high priority, but with nobody in charge. Always pick a destination.

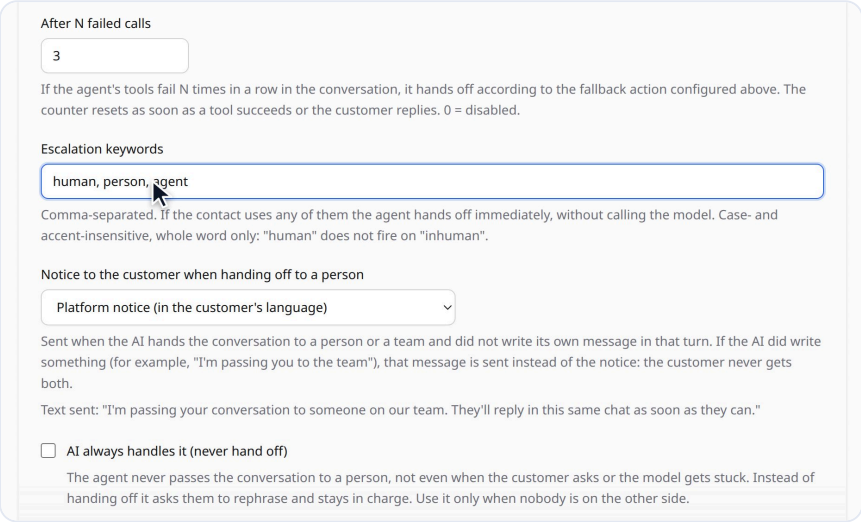
> **NOTE**
>
> With a destination chosen, escalation applies even if the “Assign to agent or team” tool is off. Turning it on is for something else: letting the AI hand off on its own when it sees fit and pick the team for each case.

After N failed calls

If the agent’s tools fail that many times in a row in the conversation, it escalates. The counter goes back to zero when a tool works or the customer replies. **0** turns it off.

Escalation keywords

Comma-separated words; if the customer writes one, the conversation is transferred immediately, without going through the AI. Case- and accent-insensitive, whole word only: “human” does not trigger on “humans”. Happy Paws uses **human, person, agent**.



After N failed calls, the keywords “human, person, agent” and the notice to the customer.

Notice to the customer when handing off to a person

Option	What the customer gets
Platform notice (in the customer’s language)	A standard text in their language, shown on screen.
My own text	Whatever you write in **Notice text** (required). E.g.: “I am passing you to the Front desk so they can help you. 🐾”
Sent by a workflow	A workflow with the “The AI handed off the conversation” trigger sends it (chapter 40). If the workflow sends nothing, the **Fallback text (optional)** goes out.
Do not send a notice	Nothing.

In every case, if the AI already wrote something at that moment (for example, “I’m passing you to the Front desk”), that is sent and not both.

AI always handles it (never hand off)

It does the opposite: the agent never passes the conversation to a person, not even when the customer asks. Instead it sends **What to say when it cannot hand off** (“Sorry, I didn’t quite catch that. Could you rephrase it?” if you leave it empty). With this on, keywords and the replies cap do not apply.

> **HEADS UP**
> Only use it if there is nobody serving behind it. If you have a team, let the AI hand off.

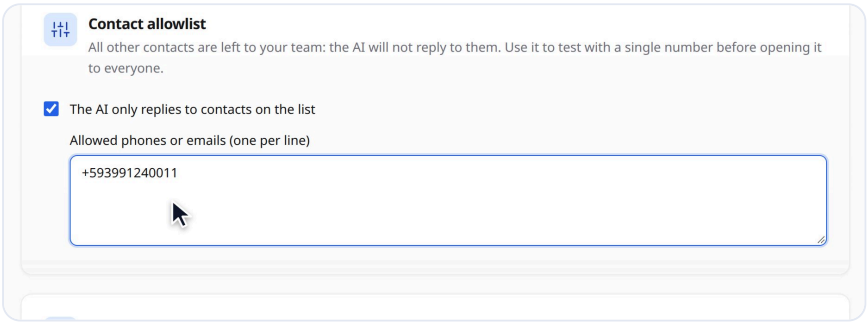
28

CHAPTER 28

Contact allowlist

Launching an agent with your own number.

The **Contact allowlist** is perfect for launching an agent on the real channel without your customers seeing it. If you turn on **The AI only replies to contacts on the list**, the AI only serves the phones or emails you write in **Allowed phones or emails (one per line)**; your team serves everyone else.



The allowlist on, with a single number.

1. Turn the list on and write your own number (for example +593991240011). Commas or semicolons also work; up to 200 entries.
2. Save and message the agent from your phone, on the real channel.
3. When you are happy, turn the list off to open the agent to everyone.

A phone number is recognized even if written differently: the last 9 digits are compared (09... matches +593...). The list also applies to inactivity follow-ups and to calls.

> **HEADS UP**
> With the list on and empty, the AI answers nobody (an amber warning shows). And a contact with no phone or email, as is common on Instagram or Messenger, never matches.

29

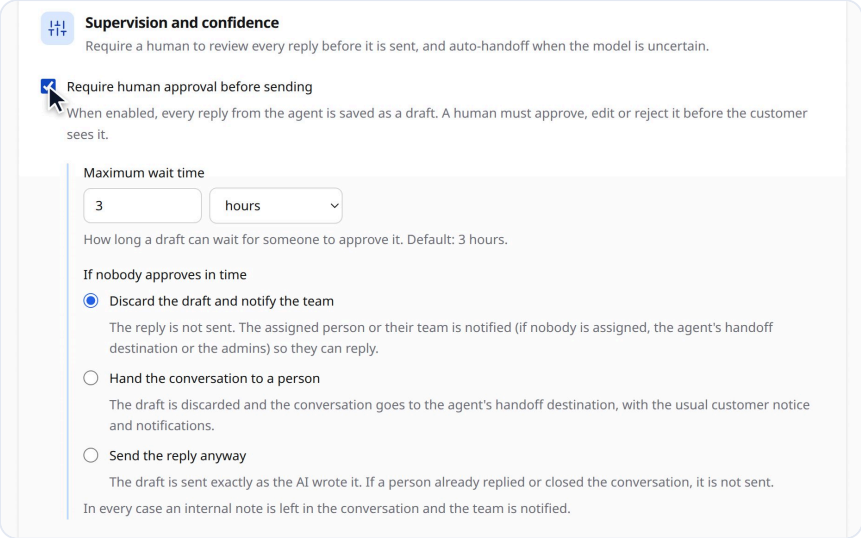
CHAPTER 29

Supervision and confidence

Human review and safety nets.

Require human approval before sending

Turns every AI reply into a **draft**: someone on your team approves, edits or rejects it from the inbox before it reaches the customer (chapter 38). Ideal for a new agent's first weeks.



Require human approval: maximum wait time and what to do if nobody approves in time.

Field	What it does
Maximum wait time	How long a draft can wait, in minutes or hours: from 1 minute to 7 days (3 hours by default).
If nobody approves in time	**Discard the draft and notify the team** (not sent; whoever is responsible is notified so they can reply), **Hand the conversation to a person** (goes to the escalation destination) or **Send the reply anyway** (unless a person already replied or closed it).

In every case an internal note is left. If the customer writes something new or a person replies first, the draft expires and nothing is run.

Confidence threshold for handoff

The AI rates how sure it is of each reply, from 0 to 1. If it falls below the value you set (for example 0.70), the conversation goes to a person. **Disable** turns it off (default). It is the AI's self-assessment: use it as a safety net, not as your only rule.

Escalate to a human after N replies without resolving

A fixed lock: after that many AI replies in the same stretch (since a person last replied), the conversation goes to your team, with a summary for the advisor. Replies to a “thanks” or an “ok” do not count. Empty or 0 turns it off; maximum 50.

30

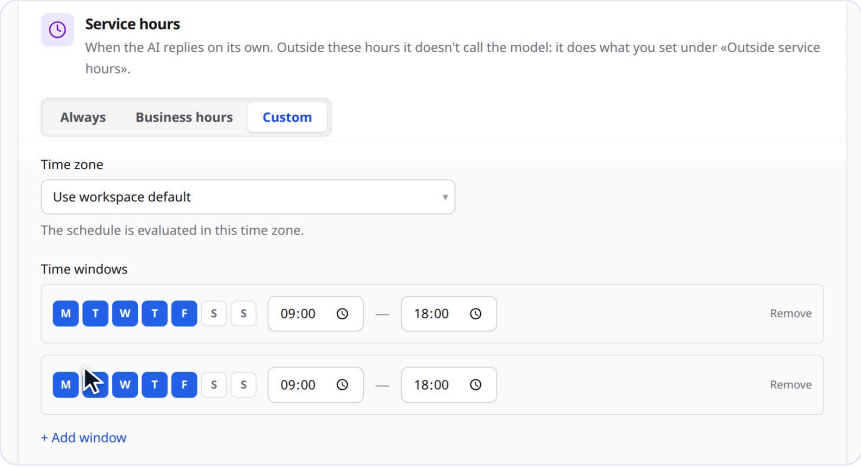
CHAPTER 30

Service hours and outside service hours

When the AI replies on its own and what it does outside those hours.

Service hours

Mode	When the AI replies
Always	At any time, every day (default).
Business hours	Your workspace’s business hours, shown day by day with their time zone and the **Edit business hours** link. If you change them there, the agent follows. If your workspace has no hours set, the AI always replies.
Custom	You choose the **Time zone** (or “Use workspace default”) and the **Time windows**: days, start time and end time. **+ Add window** adds another window.



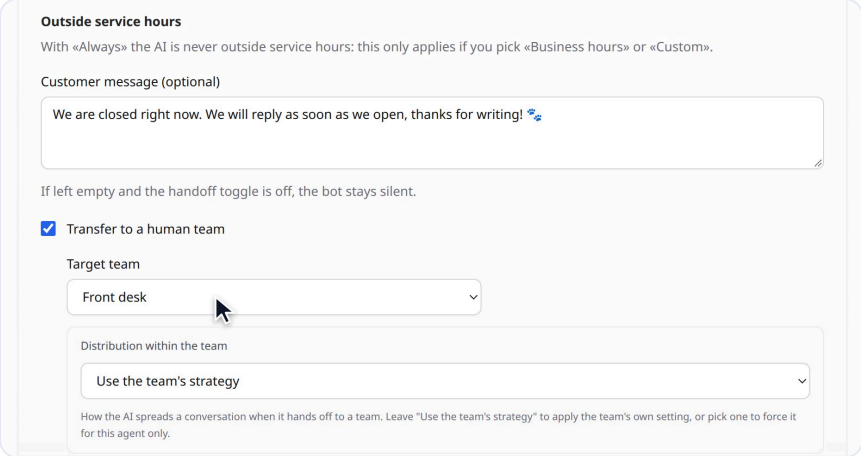
Custom hours: time zone and two time windows.

For example: Monday to Friday from 09:00 to 19:00 and Saturdays from 09:00 to 14:00, in two windows.

> **TIP**
> The end time must be after the start time. To cover the night (for example 22:00 to 06:00), use two windows: 22:00–23:59 and 00:00–06:00.

Outside service hours

Outside service hours the AI does not call the model: it does what you set here.



Outside service hours: the customer message and the transfer to the Front desk team.

Field	What it does
Customer message (optional)	A notice, for example: “We are closed right now. We will reply as soon as we open, thanks for writing! 🐾”. It is sent only once per out-of-hours stretch, even if the customer writes several times.
Transfer to a human team	Passes the conversation to a **Target team** (for example, whoever is on call), with its **Distribution within the team**. With no team, it stays unassigned in the inbox.

If you leave the message empty and the transfer off, the AI says nothing outside service hours. With the “Always” mode this section is not used.

31

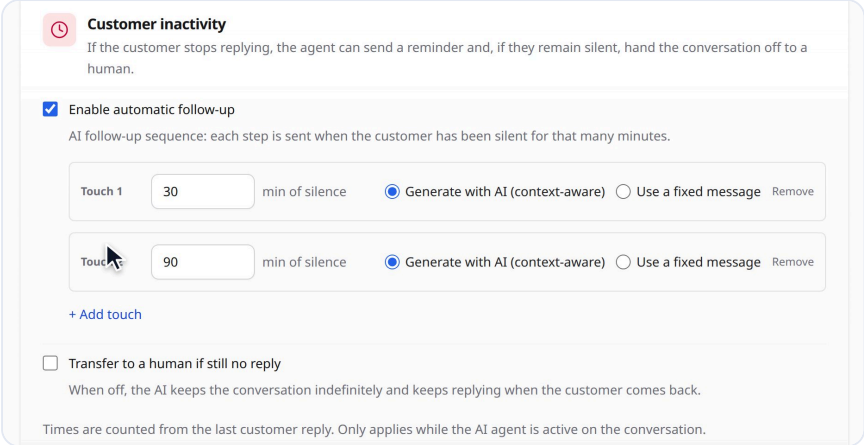
CHAPTER 31

Customer inactivity

When the customer stops replying.

Automatic follow-up

Turn on **Enable automatic follow-up** and build a sequence of touches: each **Touch** is sent when the customer has been silent for that many minutes.

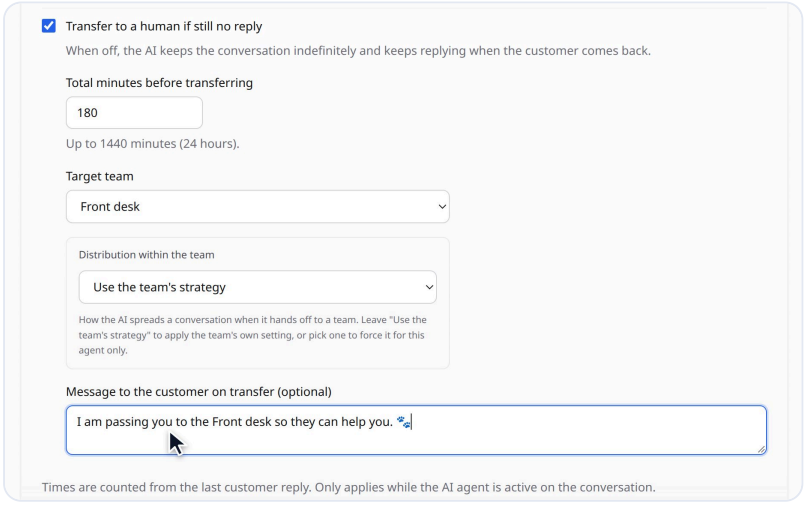


Automatic follow-up: two AI-generated touches.

Part of the touch	What it is
Minutes	"N min of silence", counted from the customer's last message. Each touch, more than the previous one.
Generate with AI (context-aware)	The AI writes the reminder based on the conversation.
Use a fixed message	Your own text, always the same. E.g.: "Are you still there? If you don't need anything else I'll close the conversation."
+ Add touch / **Remove**	Up to 10 touches.

Transfer if there is still no reply

With **Transfer to a human if still no reply** the conversation goes to a **Target team** after the **Total minutes before transferring** (up to 1,440, that is 24 hours, and always after the last touch), with its distribution and a **Message to the customer on transfer (optional)**.

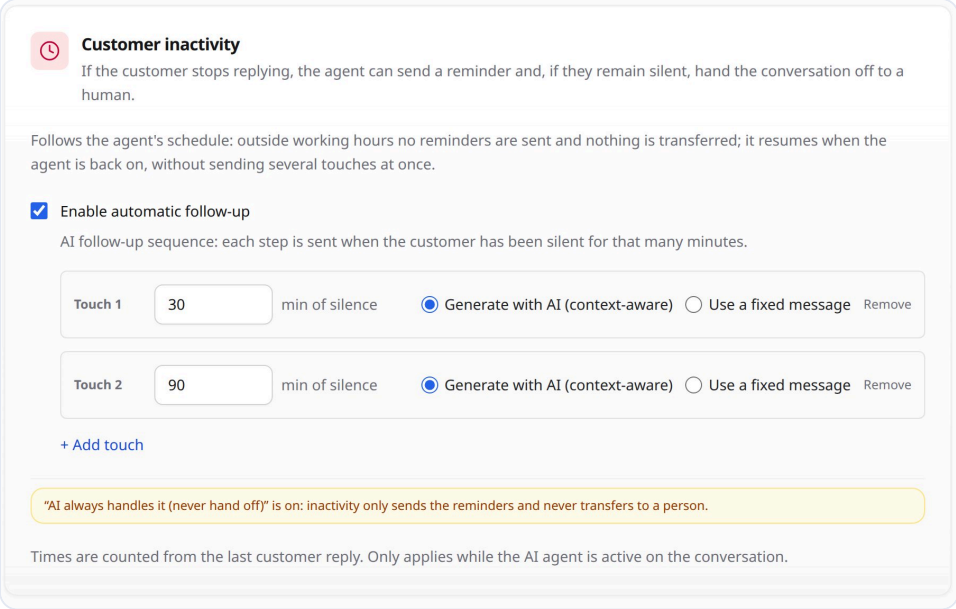


Transfer to a human if still no reply: total minutes, team, distribution and message.

If you leave it off, the AI keeps the conversation and replies when the customer comes back.

The agent's schedule and “never hand off”

The inactivity follow-up respects two settings from other sections, and says so inside this same section:



Inactivity with a schedule and “never hand off”: the two new notices and, without the transfer section, only the reminders.

If…	What inactivity does
The agent has **Working hours** (chapter 30)	Outside working hours it sends no reminders and transfers nothing: it waits until the agent is back on and continues the sequence where it was. When it reopens it keeps the spacing between touches (it does not send several at once), and the transfer also waits its time after the last touch. Notice: “Follows the agent's schedule...”.
AI always handles it (never hand off) is on (chapter 27)	It only sends the reminders: it never transfers to a person because of inactivity. The “Transfer to a human if still no reply” part is hidden (what you had set up is not deleted) and an orange notice appears. Also, the reminder the AI writes does not offer or promise that a person will take over.

An example: working hours Monday to Friday from 9:00 to 18:00 and touches at 30 and 90 minutes. If a customer stops replying on a Tuesday at 17:50, the first touch does not go out at 18:20 (already closed): it goes out on Wednesday at 9:00, and the second one at least an hour later.

> **NOTE**
>
> A sequence does not start if the customer has been silent for more than 24 hours; nor when coming back from outside working hours, if by then more than 24 hours have passed. On WhatsApp, remember you can only write freely within 24 hours of their last message.

32

CHAPTER 32

Memory, summary, style and anti-spam

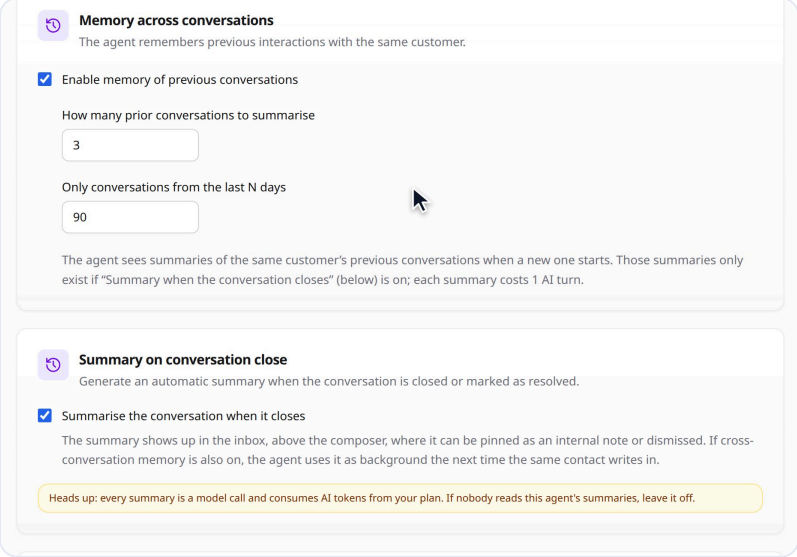
Remembering the customer, summarizing, matching the tone and not replying in pieces.

Summary on conversation close

Summarise the conversation when it closes generates an automatic summary when the conversation is closed or resolved. Your team sees it in the inbox (chapter 38). Each summary is an AI call and uses tokens: turn it on if someone will read them, or if you use memory.

Memory across conversations

Enable memory of previous conversations: when the same customer comes back, the agent sees the summaries of their previous conversations and "remembers" what happened (their dog's name, the pending vaccine).



Memory across conversations and Summary on conversation close, both on.

Field	Range
How many prior conversations to summarise	From 1 to 10 (3 by default).
Only conversations from the last N days	From 1 to 365 (90 by default).

HEADS UP

Memory uses the summaries: for it to have memory, "Summarise the conversation when it closes" must be on.

Reply style

Option	What it does
Fixed tone	Always keeps the style of your instructions (default).
Adaptive tone	Matches how the customer writes: more formal or more casual, with or without emojis.

Anti-spam

Batch consecutive messages (on by default): if the customer sends several short messages in a row (“hi”, “a question”, “how much is a bath?”), the agent waits and replies once with the full context.

Field	What it is
Wait (seconds) before replying	From 1 to 60 (5 by default).
Maximum messages to batch	From 1 to 20 (5 by default). If more arrive, the agent still replies only once.

> **NOTE**
> On WhatsApp, Instagram, Messenger, Telegram, TikTok and SMS the real wait is at least 15 seconds (20 if the customer keeps typing), even if you set less. Lower values only make a difference on the website chat and email.

CHAPTER 33

33 Take over on assignment and quality scoring

Replying on assignment and scoring conversations.

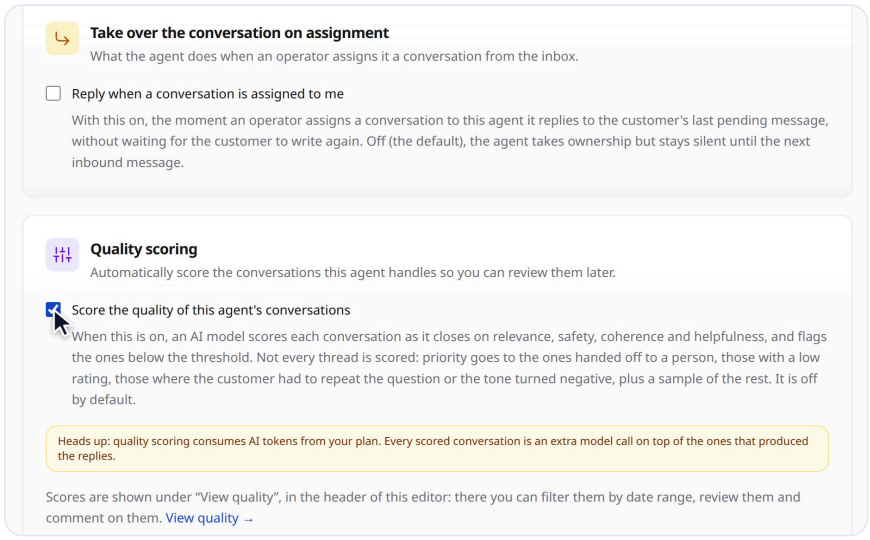
Reply when a conversation is assigned to me

When someone on your team assigns a chat to the agent from the inbox, the agent answers the customer’s last pending message right then, without waiting for them to write again. It starts off.

When it does not reply	Why
WhatsApp with the 24-hour window expired	Meta only lets you message freely within 24 h of the customer’s last message. An internal note is left.
Bulk assignments	So it does not answer dozens of chats at once.
Resolved or closed conversations, or outside the agent’s hours	—
Reassigning the same conversation to the same agent	It replies only once per message.

Quality scoring

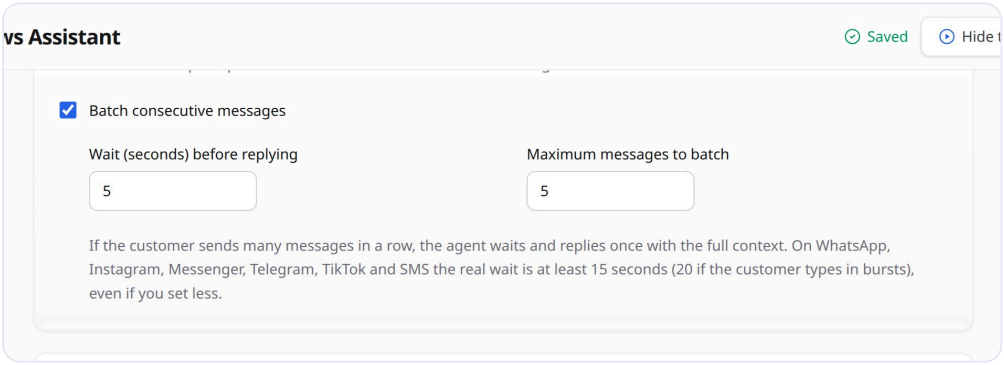
Score the quality of this agent's conversations: when a conversation closes, another AI scores it on relevance, safety, coherence and helpfulness, and flags the ones worth reviewing. Conversations handed to a person, with a low rating, where the customer repeated the question or with negative sentiment are always scored, plus a sample of the rest. Scores are shown under **View quality** (chapter 42).



Take over on assignment, and Quality scoring on with its usage warning.

> **HEADS UP**
> Scoring uses tokens from your plan: every scored conversation is one more AI call. It starts off.

We set up what the vet clinic needs and click **Save**.

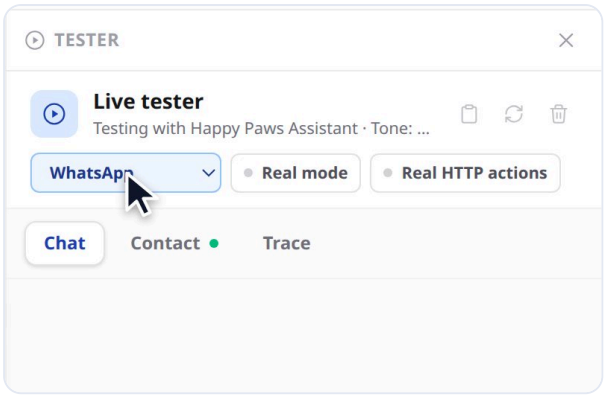


Anti-spam with its default values, and the agent saved.

34 CHAPTER 34 — The tester: chat, channel and contact

Testing before any customer sees it.

The **Live tester** sits to the right of the editor (**Test AI agent** shows it). It uses the same engine as the real channel, with the agent's **saved** version: its instructions, its knowledge base, its tools and its level. At the top it says "Testing with … · Tone: …".



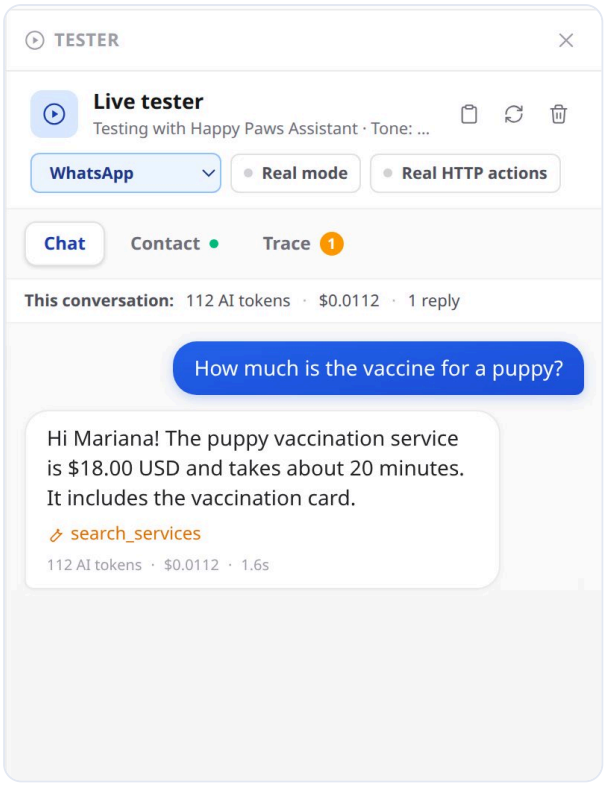
The tester with WhatsApp as the simulated channel.

The header

Control	What it does
Simulated channel	"Channel: not set", WhatsApp, Webchat, Telegram, Instagram, Messenger, Email, SMS or Voice call. The agent replies as if it were on that channel (for example, with no formatting on WhatsApp) and only uses that channel's tools.
Copy	**Copy conversation (JSON)**, to share the test with support.
Regenerate	**Regenerate the last reply**.
Restart	**Restart conversation**: start over.
Real mode and **Real HTTP actions**	See chapter 35.

The chat

Write as a customer would and click **Send** (or Enter). Under each reply you see the AI tokens, the cost and the time; in amber, with a wrench, the tools it used. At the top the usage of the whole test adds up ("This conversation: ... AI tokens · ... replies").



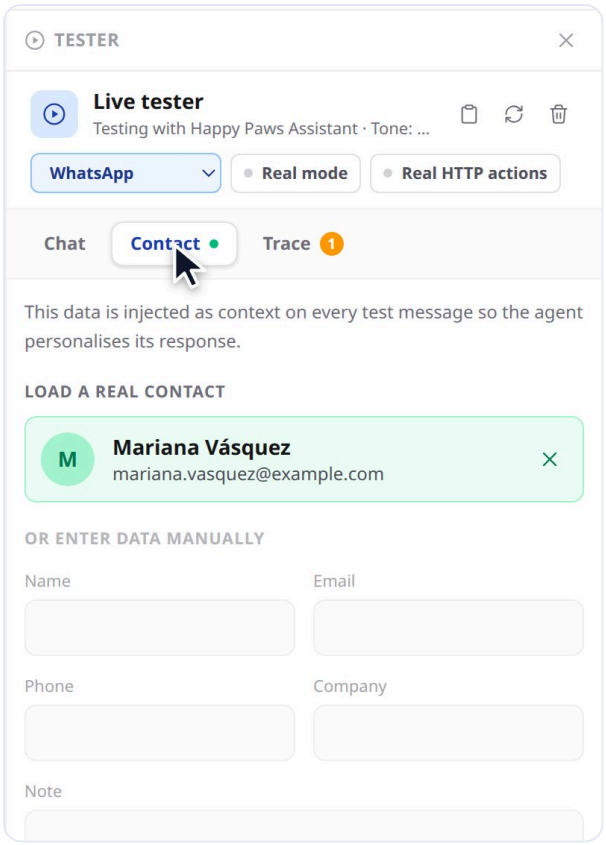
A reply in the tester, with the tool used and the tokens, cost and time below.

When the AI decides to pass the conversation to a person, a grey bubble appears: “In the real channel, the conversation would be handed to a person here”, with the destination.

> **HEADS UP**
>
> The tester does use AI balance, just like a real conversation. You can send up to 30 messages a minute.

The Contact tab

“This data is injected as context on every test message so the agent personalises its response”. You can **Load a real contact** (search by name, email or phone) or, **Or enter data manually**: Name, Email, Phone, Company, Note and Tags. That way the agent addresses the “customer” by name and you can test the variables.



The Contact tab: load a real contact or type the details by hand.

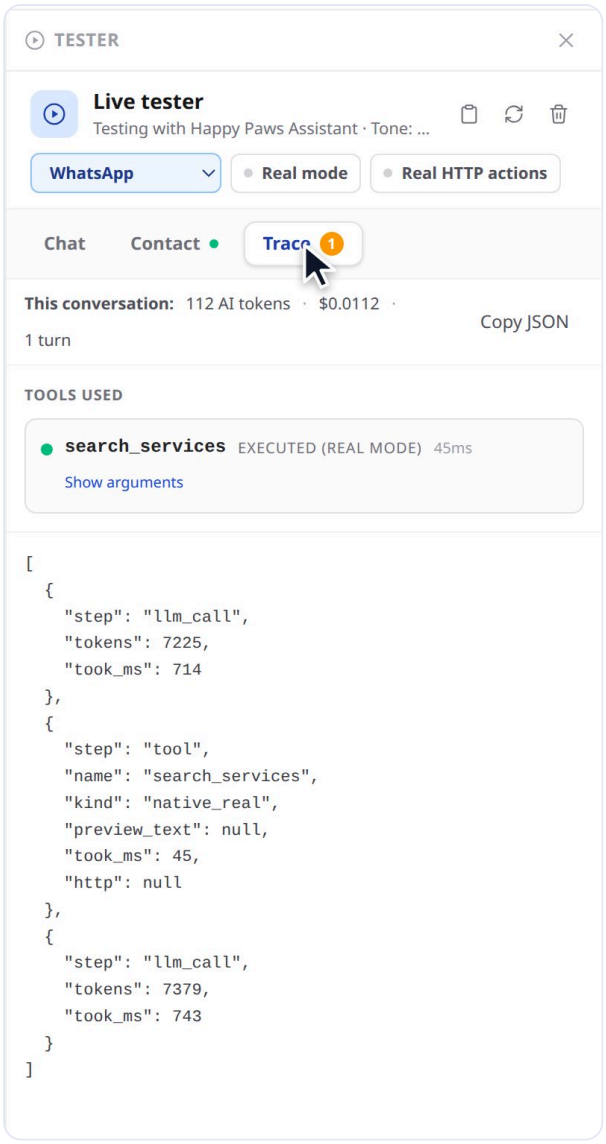
CHAPTER 35

35 The tester: trace, real mode and more

See why it answered what it answered, and test for real.

The Trace tab

It shows the **Tools used** in the last reply: which one, with what data (**Show arguments**), what it returned and how long it took, with a colored dot: green if it really ran, amber if it was simulated, red if there was an error. Below, the full detail, with **Copy JSON**. It is the best way to understand a reply.

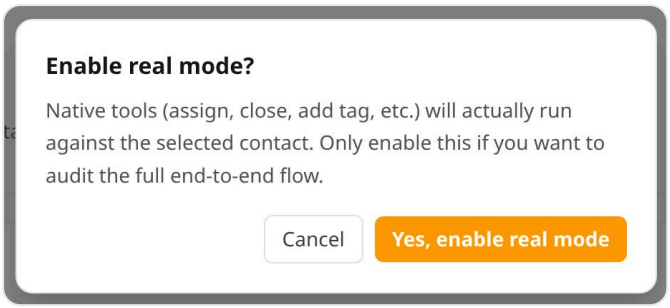


The Trace: the tools used in the last reply, with their arguments and result.

Simulated or real

Kind of tool	With Real mode off (default)	With Real mode on
Lookups (search services, check times, search the catalog)	They really run: you see real data.	Same.
Ones that change data (book, assign, close, tag, create orders)	**They are simulated**: the trace says "Simulated" and shows what would have happened. Nothing is created.	They really run on the loaded contact: the appointment is created, the conversation is assigned.

Real mode can only be turned on with a **real contact** loaded in the Contact tab, and the first time it asks you to confirm ("Enable real mode?"). When on, a red strip warns "REAL MODE — actions modify workspace data". Reply texts are not sent to the customer: only the tools' effects happen.



Enable real mode?: the warning before really running tools.

Real HTTP actions does the same for your custom actions: off, the trace shows “NOT sent. This is what would have been sent”; on, it really calls your system (“Actually SENT to your endpoint”). Sensitive values (keys, tokens) are hidden in the trace.

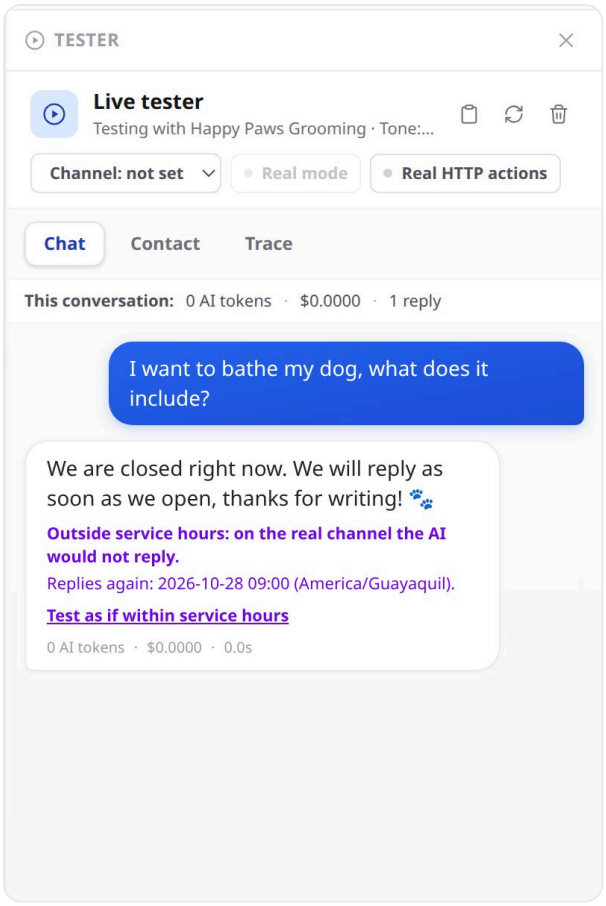
> **HEADS UP**
> Only use real mode to test the full flow, with a test contact (for example, yourself). If that contact has an open conversation, the tools act on it.

Attachments and voice notes

With the paperclip (**Attach audio, image or PDF (uses credits)**) you send an image, an audio file or a document, up to 25 MB, and with the microphone (**Record voice note**) you record audio; that is how you test what the agent understands. If the matching capability is off in Tools, the tester tells you instead of processing it.

Outside service hours

If you test outside the agent’s hours, the tester does not call the AI and shows what would happen on the real channel: “Outside service hours: on the real channel the AI would not reply”, the out-of-hours message, which team it would go to and when it replies again. It costs nothing. The **Test as if within service hours** button ignores the hours during the test (**Respect service hours** applies them again).



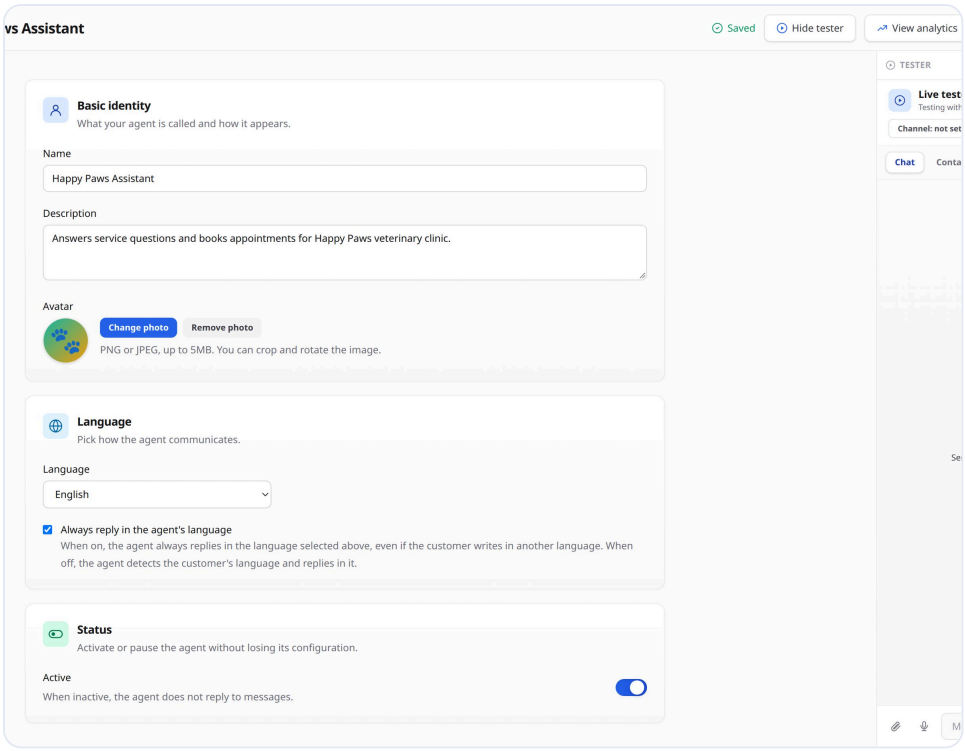
Testing outside service hours: what would happen on the real channel and the Test as if within service hours button.

CHAPTER 36

36 Activating the agent and a real example

Switching it on, and watching it serve a real customer.

When you are happy, click **Save and activate**. The agent goes online and starts answering on the connected channels: the **Active** switch stays on, the indicator says **Saved** and on the list the agent shows as **Online**.



Saved and active: the Active switch on and, at the top, Saved.

An appointment from the website chat

Carla Ríos opens the Happy Paws website chat:

Hi, how much is a bath for my dog?

Hi! The bath and grooming service is $20.00 and takes 1 hour. If you like, I can also help you book it for your dog.

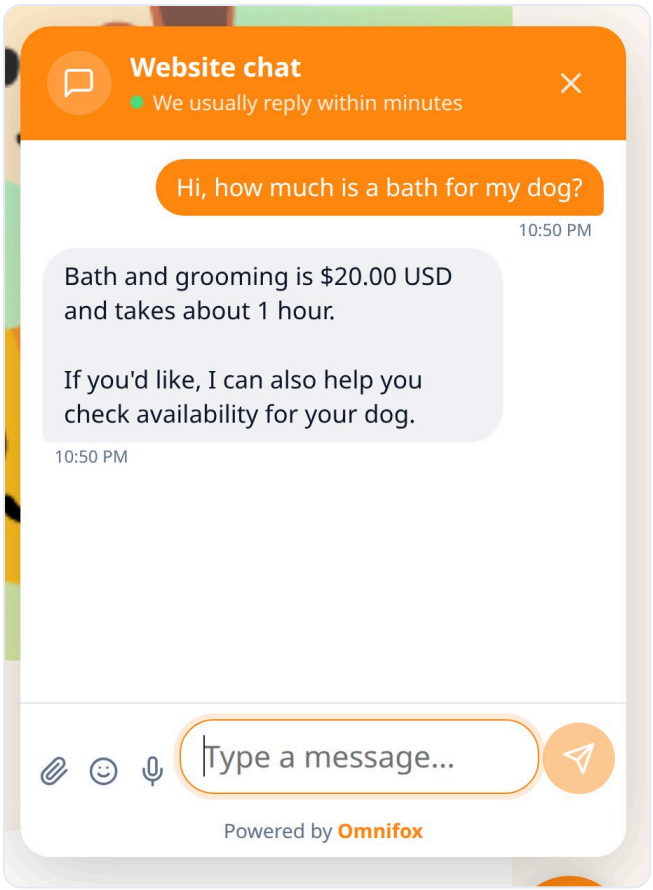
Yes, I would like a bath for my dog Rocky on Saturday. I am in Quito

On Saturday I have 09:00, 10:00 and 11:00 free. Which one works for you?

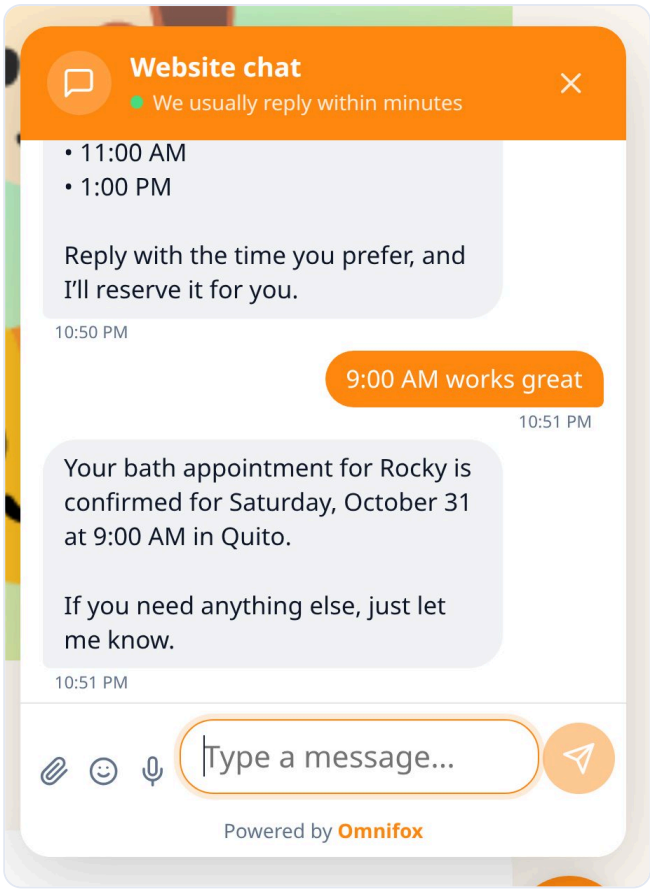
09:00 works great

Done! Rocky's bath is booked for Saturday, October 31 at 09:00. 🐶

The agent looked up the service (Search services) and answered with the real price; it checked the calendar (Check appointment availability) and offered truly free times; and when Carla picked one, it booked the appointment (Book appointment). All in seconds, without anyone on your team stepping in.



The Happy Paws website chat: the agent answers the real price of the bath.



The appointment booked from the chat, with nobody on the team stepping in.

NOTE

The times the agent offers come from the real Calendar. The appointment shows in the calendar with its origin: booked by the AI agent.

CHAPTER 37

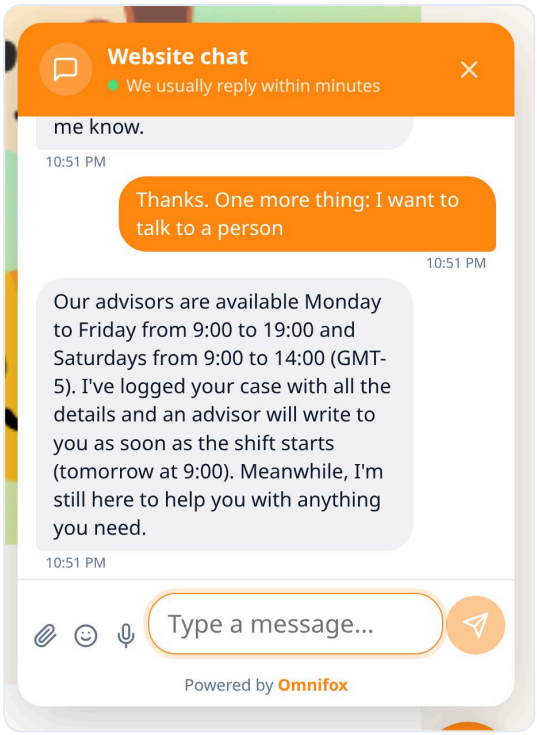
37 When the customer asks for a person

A keyword and the conversation goes to the Front desk.

Thanks. One more thing: I want to talk to a person

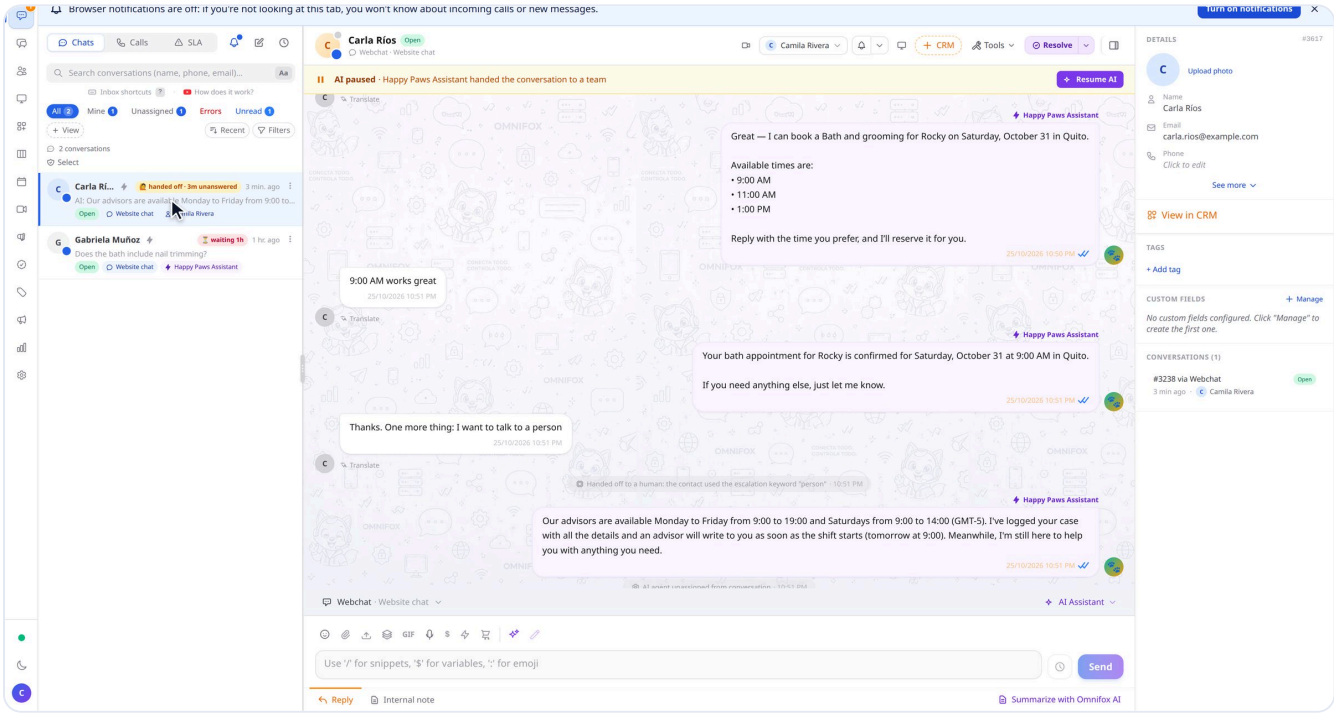
Handed off to a person: the contact used the escalation keyword “person”

“Person” is one of its keywords, so the agent passes the conversation to the Front desk team immediately, without asking the AI, and tells the customer. If your team is outside its hours at that moment, the notice says so: an advisor will write as soon as the hours start.



The customer asks for a person: the conversation goes to the Front desk and the agent tells her when they will reply.

In the inbox, the Front desk sees the whole conversation: what the customer wrote and what the AI answered, with the agent's mark.



The conversation in the inbox: what the customer wrote, what the AI answered and the hand-off note.

38

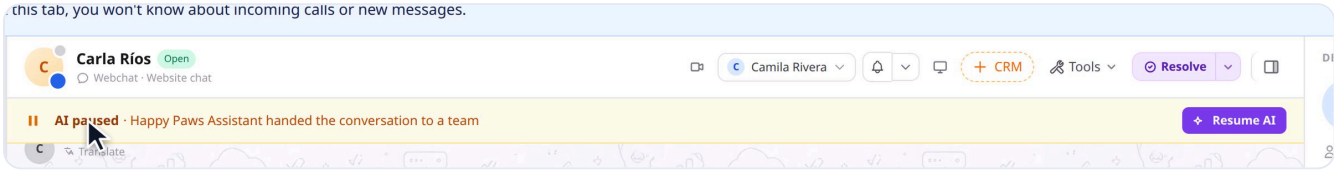
CHAPTER 38

The AI in the inbox

How the AI lives alongside your team in the inbox.

AI paused

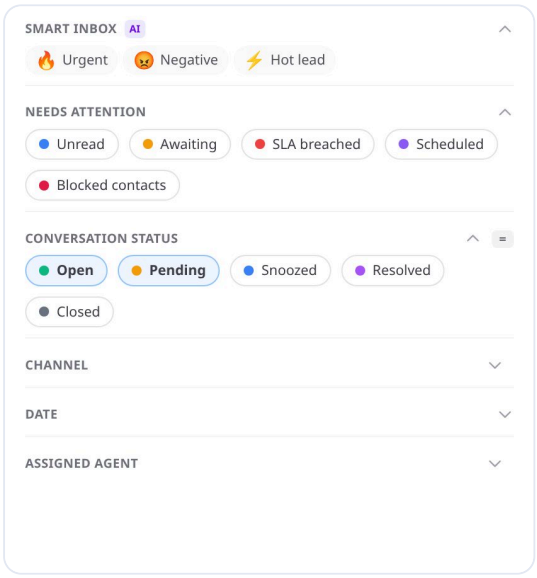
When the AI leaves a conversation, **AI paused** shows at the top with the reason: you took over, it handed the conversation to a team, it moved to an after-hours team, it handed it to you after the customer went quiet, or someone replied from the WhatsApp app. From then on your team works as usual and the AI stays out of it.



AI paused: "Happy Paws Assistant handed the conversation to a team", with the Resume AI button.

If a person replies in a chat the AI is handling, they take over and the AI pauses. To give it back, use **Resume AI**, or assign the conversation to an AI agent with the **Assign** button (AI agents section). You cannot assign it to an agent that is off: Omnifox warns you it would not reply.

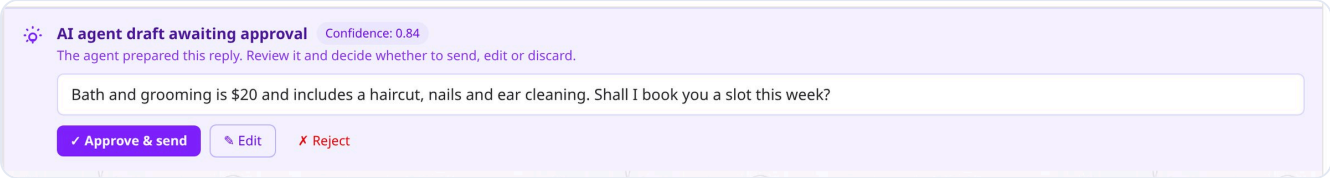
In the inbox filters you can show only one AI agent's conversations.



The inbox filters; under Agent Assigned you can pick an AI agent.

Drafts to approve

With human approval on (chapter 29), the AI's replies arrive as **AI agent draft awaiting approval**, with the text it prepared and its confidence level.



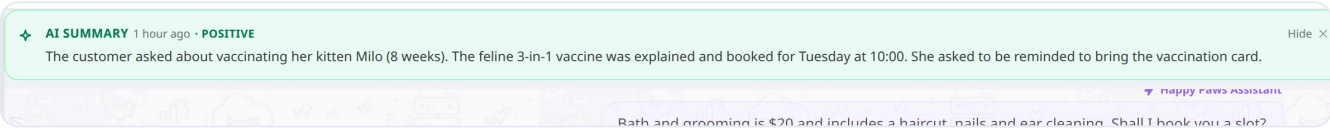
An agent draft awaiting approval: Approve & send, Edit or Reject.

Button	What it does
✓ **Approve & send**	The draft goes out as is.
✎ **Edit**	You correct it and click **Send edit**.
✗ **Reject**	It is not sent; you can note a **Reject reason (optional)**.

The customer only receives what a person approved.

The conversation summary

With the summary on close turned on (chapter 32), every closed conversation leaves an **AI summary** above the composer, with the customer's sentiment.



The AI summary of a closed conversation, with the customer's sentiment.

From the menu, **Summarize conversation** opens the summary with what was left pending and the suggested actions; you can copy it or pin it as an internal note.

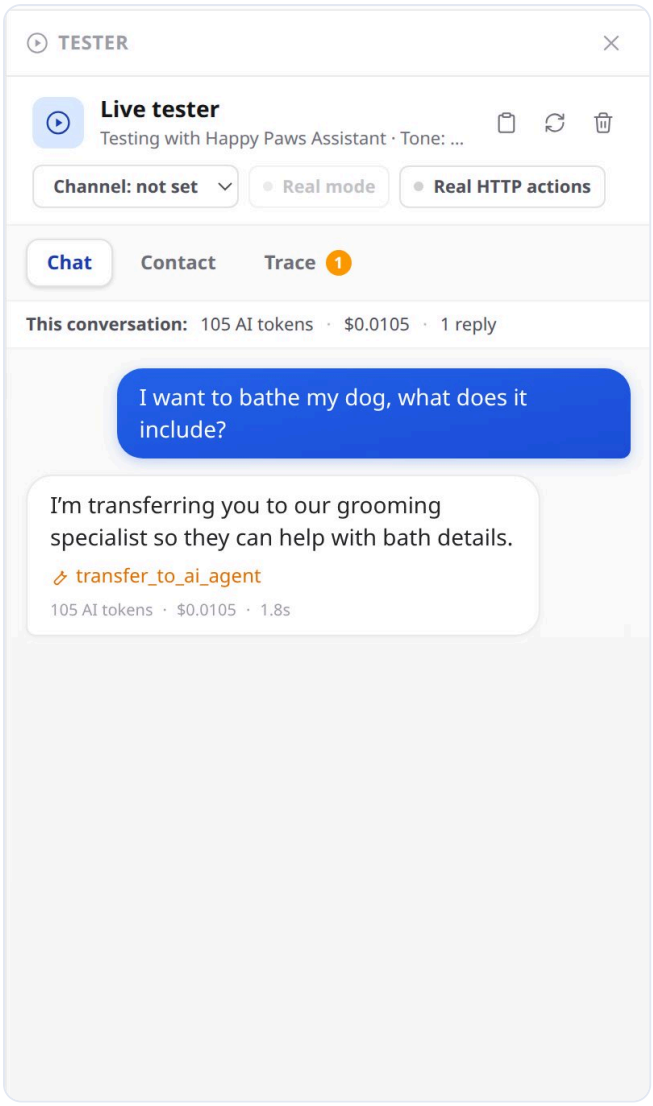
39

CHAPTER 39

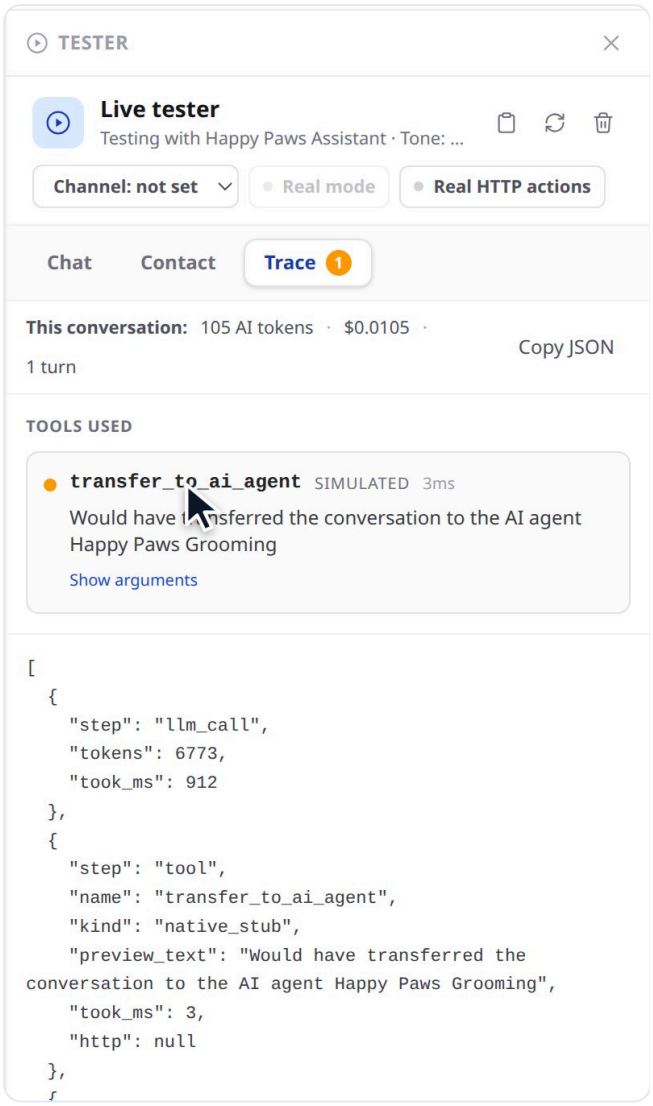
Several agents working together

One agent per specialty, passing the conversation to each other.

You can have several specialized agents. Happy Paws has a second agent, **Happy Paws Grooming**, an expert in baths and haircuts. With the **Transfer to another AI agent** tool on in the front desk agent, when the customer asks about grooming, the front desk passes the conversation over. Its instructions say: "Bath and grooming: do NOT answer yourself. Transfer the conversation to Happy Paws Grooming right away with !transfer_to_ai_agent".



The customer asks about the bath and the front desk uses Transfer to another AI agent.



The trace: "Would have transferred the conversation to the AI agent Happy Paws Grooming".

In the tester the transfer is simulated; in the trace you see which agent the conversation would have gone to. On a real channel, Happy Paws Grooming takes over the chat and keeps answering.

> **TIP**
> To chain agents, give them clear, unique names and tell each one, in its instructions, when and to whom to transfer: "If they ask about bath or grooming, transfer to Happy Paws Grooming". A chat can jump between AI agents at most twice.

40

CHAPTER 40

AI agents in workflows

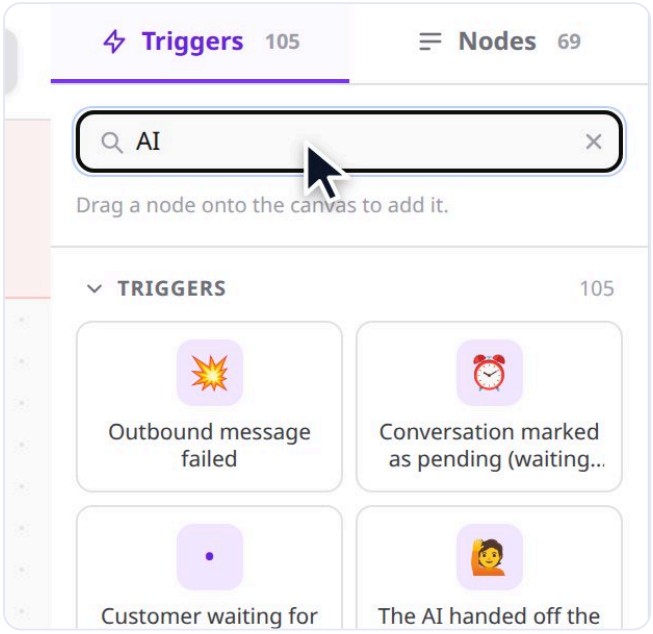
Triggers and nodes to work with AI.

In **Workflows** there is a trigger and several nodes made for AI agents.

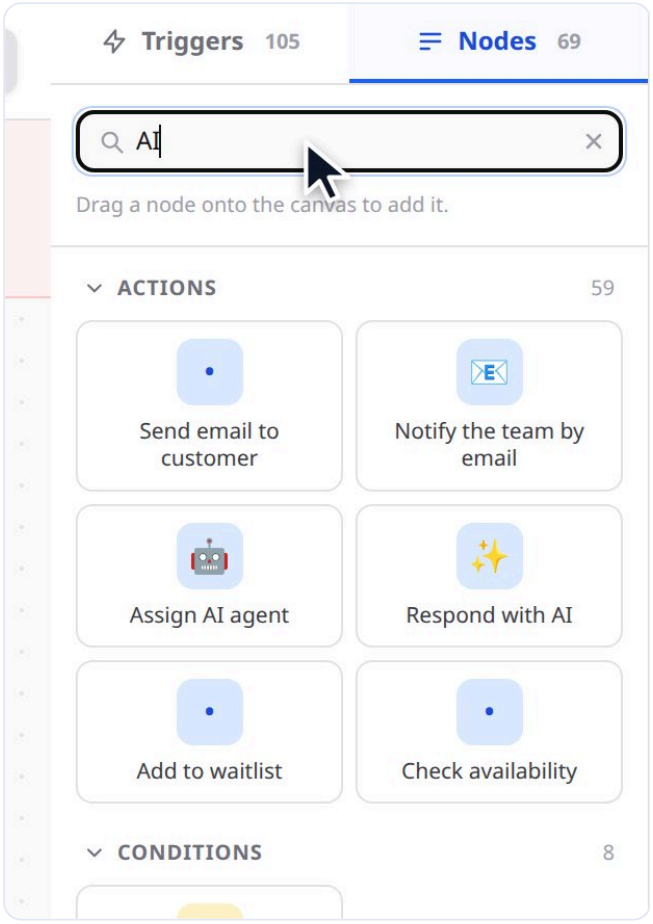
The “The AI handed off the conversation” trigger

It fires every time the AI hands off a chat: to alert a supervisor, send your own message to the customer or tag the conversation. You can filter it by **Handoff reason** (the AI assigned with its tool, keyword, turn limit, the AI asked to pass to a person, it did not reply, low confidence, failed tools, the assignment failed, plan limit reached or draft not approved in time), by destination (a person, a team or unassigned) and by team.

AI nodes



The “The AI handed off the conversation” trigger in a workflow.



The AI nodes in the workflow editor.

Node	What it does
Assign AI agent	The workflow hands the conversation to an active agent. Options: **Reply to the last unanswered message** and **Stop other workflows in this conversation**. It does nothing if a person already has the conversation or the AI is paused.
AI respond	Generates a reply inside the workflow, automatically or with your own instruction; it can send it to the channel or store it in a variable. From the Escala plan.
AI classifier (route by intent)	Picks the path based on the customer’s intent, among the categories you define (at least two). It does not talk to the customer and uses credit.

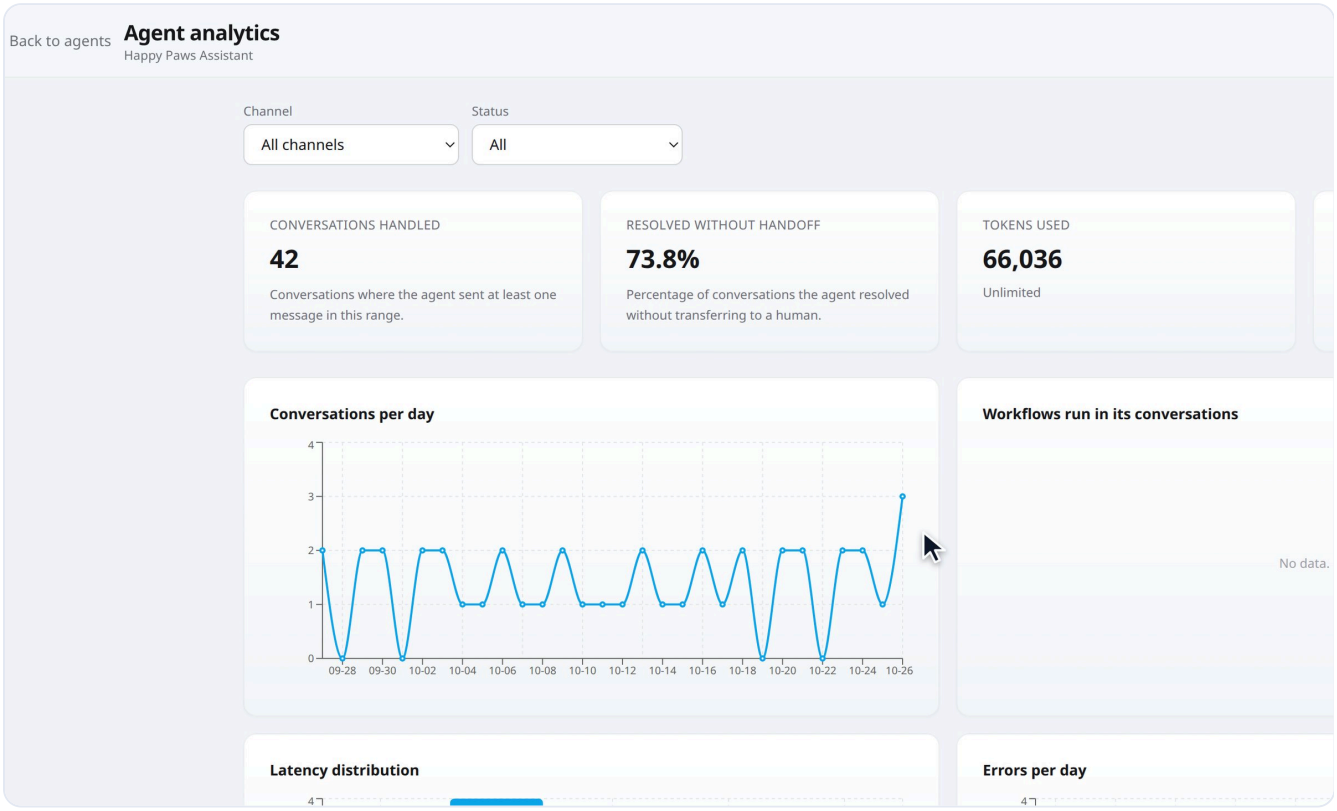
41

CHAPTER 41

Agent analytics

The agent's numbers.

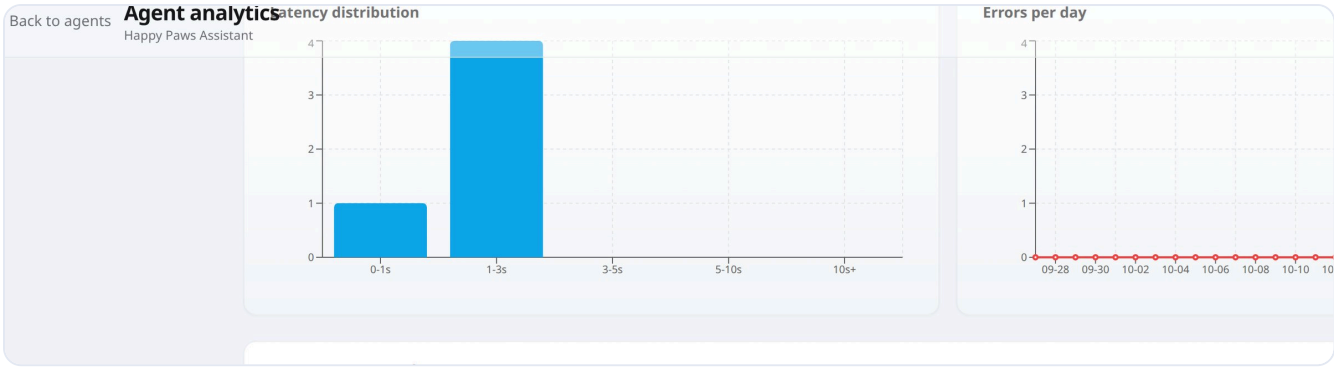
View analytics, in the editor bar, opens **Agent analytics**: its numbers for the **Last 7 days**, **Last 30 days** (default) or **Last 90 days**. You can filter by **Channel** and by **Status** (All, Resolved, Handed off).



Agent analytics: period, filters and the four cards.

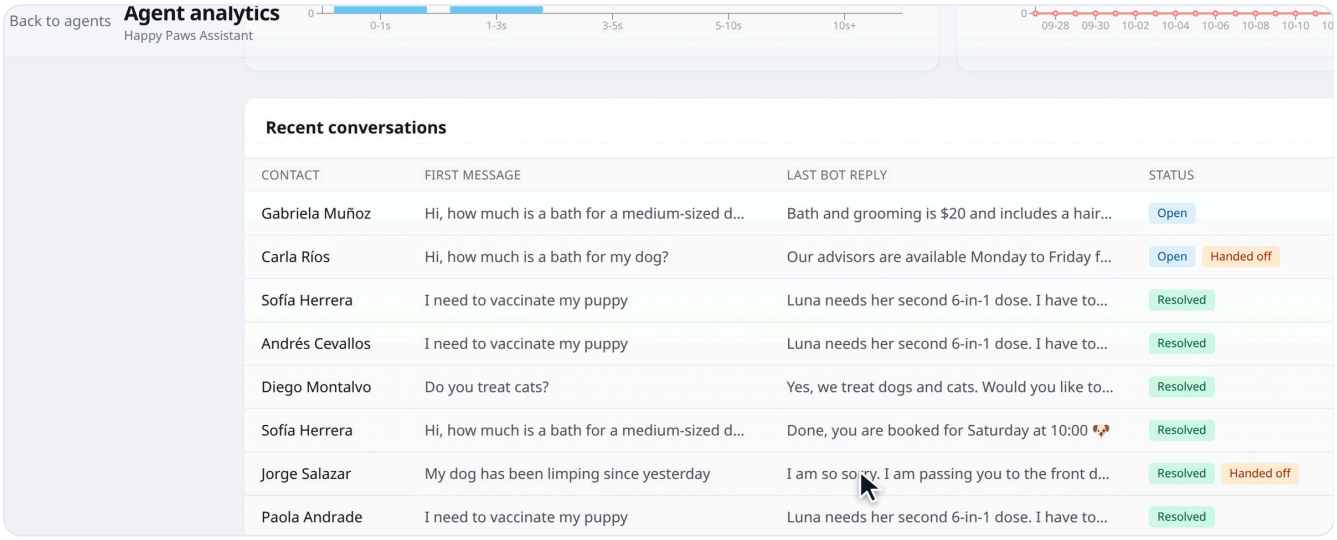
Card	What it measures
Conversations handled	Conversations where the agent sent at least one message in the range.
Resolved without handoff	Share of those conversations that were resolved or closed without going to a person.
Tokens used	Tokens the agent used in its conversations (not counting the tester).
Median latency	How long it takes to reply; next to it, the p95 latency (95% of replies take less than that).

Below, the charts **Conversations per day**, **Latency distribution**, **Errors per day** and **Workflows run in its conversations**.



The conversation, latency, error and workflow charts.

And the **Recent conversations** table: contact, first message, last bot reply, status (with “Handed off” if it went to a person), date and the **Open in inbox** link.



Recent conversations, with Open in inbox.

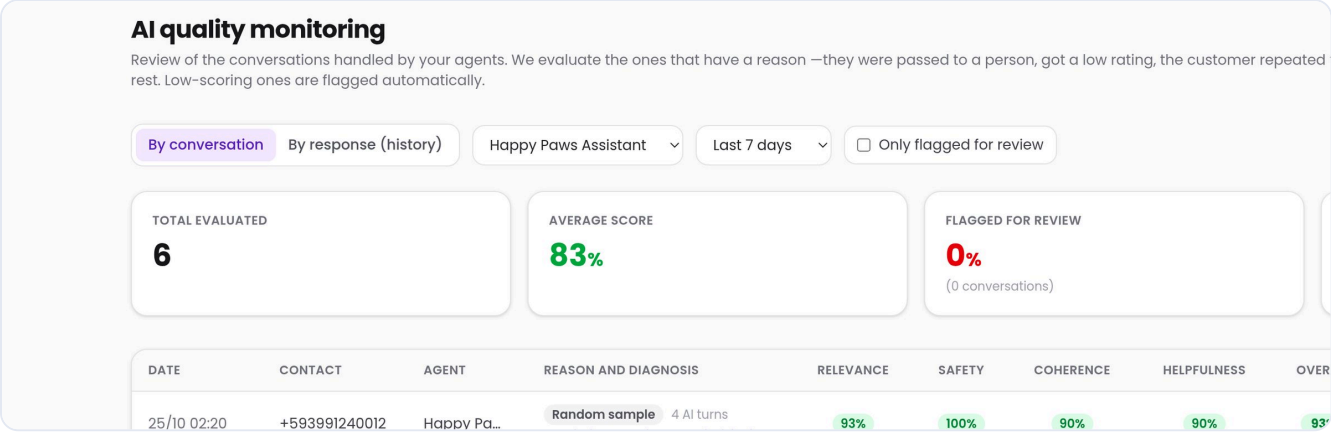
42

CHAPTER 42

Reply quality

Improving the instructions with real cases.

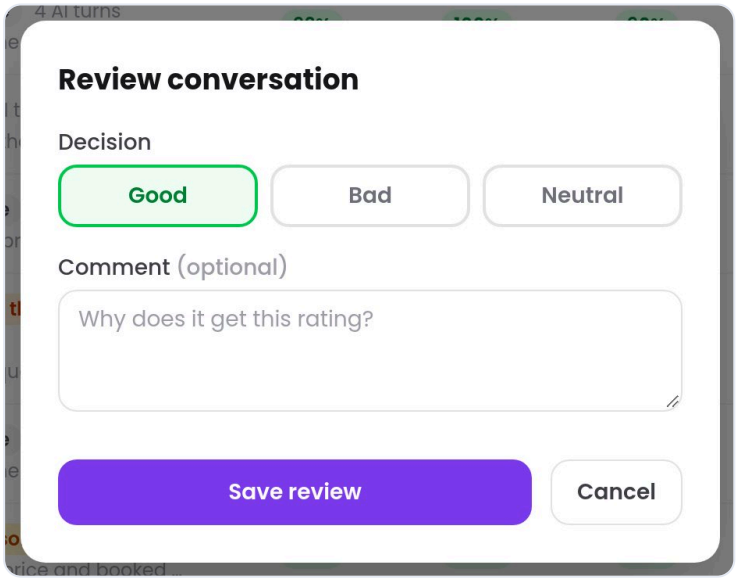
View quality, in the editor bar, opens **AI quality monitoring** with the scored conversations (it needs scoring turned on, chapter 33). If no agent has it on, the page says so and tells you where to turn it on.



AI quality monitoring: filters, cards and the scored conversations.

Part	What is there
Filters	**By conversation** or **By response (history)**, **All agents** or one, the period (**Last day**, **Last 7 days**, **Last 30 days**, **All history**) and **Only flagged for review**.
Cards	**Total evaluated**, **Average score**, **Flagged for review** and **Reviewed**.
Table	Date, contact, agent, **Reason and diagnosis** (Passed to a person, Low rating, Asked the same thing twice, Negative sentiment or Random sample), the Relevance, Safety, Coherence, Helpfulness and Overall scores, the status and the actions.

On each row you can mark it **Good** or **Bad**, or **Review** it with a decision and a comment (“Why does it get this rating?”).



Review: the decision and a comment.

TIP

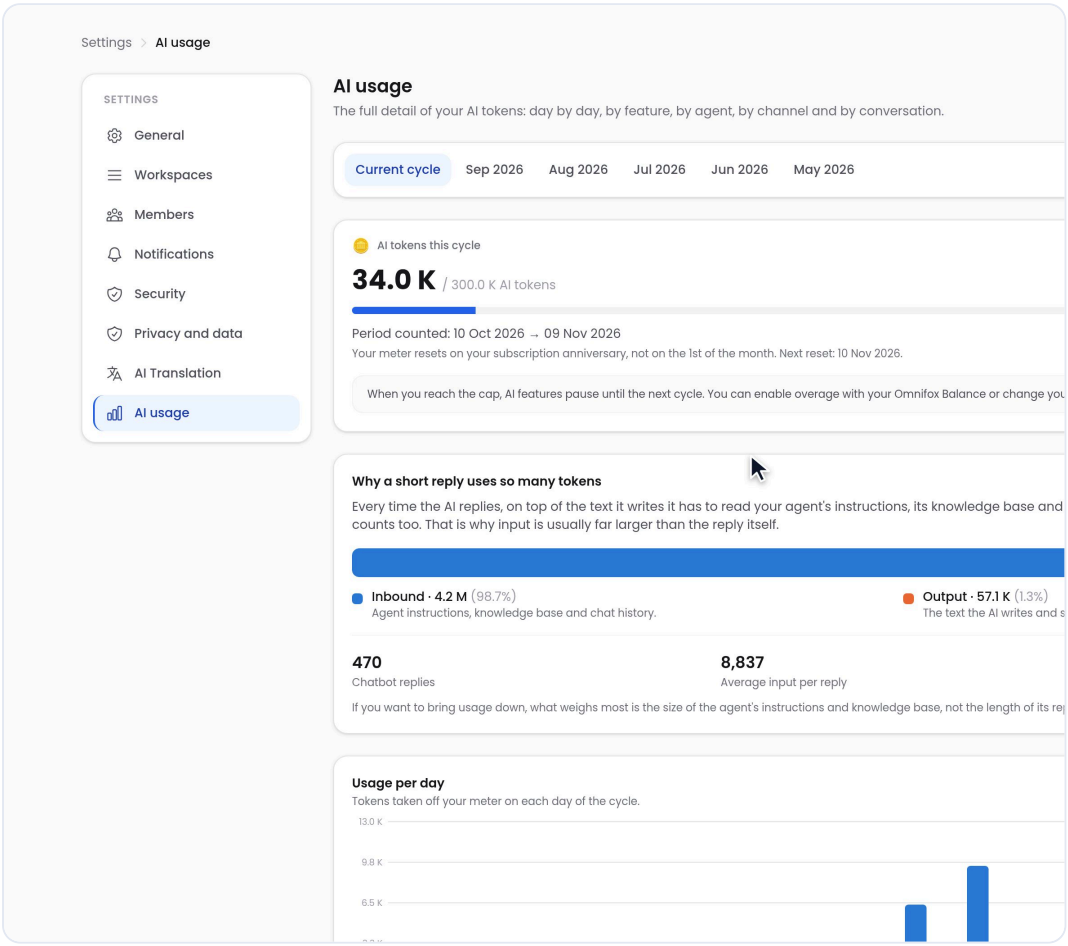
Check the flagged ones every week: each one is a clue to what is missing or extra in your instructions or your knowledge base.

CHAPTER 43

43 AI usage and credits

Everything the AI does uses tokens from your plan.

In your settings, **AI usage** shows the detail of your AI tokens: the cycle meter (with “Cap reached”, “Close to the cap” or “Unlimited”), the period counted, when it resets and the **Usage per day**. The meter resets on your subscription anniversary, not on the 1st of the month, and unused tokens do not carry over.



AI usage: the cycle meter and usage per day.

Why a short reply uses so many tokens

On every reply the AI reads its instructions, the knowledge base passages and the conversation history, and then writes. The card shows how much is input and how much is output. That is why what saves the most is having clear instructions with no filler, and turning on the expensive capabilities (photos, audio, quality, summaries) only if you need them.

What it was used for

You can see usage **By feature**, **Per agent**, **By Channel** and **By conversation** (each row opens the conversation in the inbox). Building an agent with “Describe your business” and optimizing instructions are free.

If you hit the cap

The AI pauses until the next cycle: the agent leaves a note in the conversation so your team carries on by hand and the owner gets an email. No notice is sent to the customer. If your plan allows it, you can turn on **overage** in Billing: the AI keeps working and the overage is taken from your Omnifox Balance; with a zero balance it pauses, it never creates debt.

44

CHAPTER 44

Permissions and plans

Who can do what, and what depends on the plan.

Task	Owner	Admin	Supervisor	Agent
See the list of agents and their metrics	✓	✓	✓	—
Create, edit, activate, duplicate and delete agents	✓	✓	—	—
Knowledge base	✓	✓	—	—
Create, start and apply A/B experiments	✓	✓	—	—
See quality and analytics	✓	✓	✓	—

There is no limit on the number of agents on any plan.

What changes with the plan

Feature	From
Text AI agents, AI appointments, CRM	Every plan with AI
Custom actions, courses	Impulsa
Update fields and tags, trigger workflows, book a meeting, voice on calls, automatic summary	Crece
Manage calls, AI respond node, A/B experiments	Escala
Max Fox and Deluxe Fox	Any paid plan (not on the free trial)

Whatever your plan does not include shows a padlock and the name of the plan that has it.

45

CHAPTER 45

Best practices and troubleshooting

So your agent works well from day one.

Best practices

	Practice
1	Start with **one agent per job** and few tools. Add them one at a time.
2	Launch it with the **allowlist** (your number) or with **human approval**.
3	Always pick an **escalation destination** and turn on **Assign to agent or team**.
4	Load your **FAQ** into the knowledge base, with the files the agent must send.
5	Ask it in the instructions to **never make things up**: look it up or hand off.
6	**Save before testing** and test like a real customer, watching the trace.
7	Start on **Swift Fox** and move up a level only if needed.
8	Check **quality** and **analytics** every week and improve the instructions with real cases.

Troubleshooting

Problem	What to check
The agent does not reply	That it is **active**; that it is **connected to the channel** with a strategy other than Disabled; that it is **within its hours**; that the **allowlist** is off (or has that number); that the **AI credit** has not run out; and that nobody on your team took over the conversation (AI paused).
It makes things up	Ask it to use its tools and its knowledge base and to hand off when it does not know. Lower the temperature. Check that the document is linked and Ready.
It does not use a tool	That it is **on and saved** in Tools (naming it does not turn it on); that your plan includes it; that its data exists (services, catalog...). Look at the trace in the tester.
The inactivity reminder did not go out	That you are within the agent's **working hours** (outside them it waits until it opens); that the customer has not been silent for more than 24 hours; that the AI is still in charge of the conversation.
It does not hand off to anyone	That "What to do on escalation" has a **destination**; that "never hand off" is off; that the team has members.
It replies in another language	Turn on **Always reply in the agent's language**.
It replies several times in a row	Check that **anti-spam** is on.
My custom actions are not used	That **Make HTTP requests** is on and saved, and the action is **Enabled**.

46

CHAPTER 46

Glossary and FAQ

Glossary

Term	What it is
AI agent	An assistant that answers your customers on its own, with instructions, tools and documents.
Instructions (prompt)	The text that tells the agent who it is, what it does and when to hand off.
Tool	An action the agent can decide to use: search, book, assign, close...
Custom action	A tool of your own that calls your system over HTTP.
Variable	A value filled in on each conversation, such as the customer's name ($contact.first_name).
Knowledge base	The workspace documents agents look up to answer.
Escalate / hand off	Pass the conversation from the AI to a person or a team.
A/B experiment	Comparing the current instructions (A) with another version (B) in real conversations to keep the one that works better.
Escalation keyword	A word that, if the customer writes it, passes the chat to a person immediately.
Channel strategy	When the AI steps in on a channel: Disabled, Always or New conversations only.
AI paused	The AI left a conversation (a person took over, it was transferred...) and does not reply until resumed.
Draft	An AI reply waiting for human approval before it is sent.
Engine level	Swift, Pro, Max or Deluxe Fox: higher level, more judgement and more usage.
Token	The unit used to measure what the AI reads and writes; your plan includes an amount per cycle.
Temperature	How much replies vary: low, more precise; high, more creative.
Tester	The editor's test chat, with the same engine as the real channel.
Trace	The detail of which tools the agent used, with what data and what they returned.
Real mode	In the tester, really running the tools that change data.
24-hour window	The period, since the customer's last message, in which WhatsApp lets you message them freely.
CSAT	The customer satisfaction survey, from 1 to 5.

Frequently asked questions

How many agents can I have?

As many as you want: there is no limit on any plan. "1 / 2" on the list means one active out of two.

Can I have two agents on the same channel?

No: each channel has a single chat agent. If you need specialists, connect one to the channel and use “Transfer to another AI agent” to pass them conversations.

I changed the instructions and the agent still behaves the same.

You probably did not save: the tester and the channels always use the saved version. Click **Save** (or Ctrl + S) and test again.

Does the tester use my balance?

Yes, like a real conversation. Tests outside service hours are free, and so are building an agent with “Describe your business” and optimizing instructions.

Can the AI book in my calendar?

Yes, with Search services, Check appointment availability, Book appointment and View, move or cancel the customer’s appointment. The Calendar must have its services and hours.

How do I stop the AI from answering my existing customers?

In the channel settings, choose **New conversations only**: conversations that already existed stay with your team.

I replied to a chat and the AI stopped answering.

That is expected: by replying you take over and the AI pauses. To give it back, use **Resume AI** or assign it to an AI agent.

What does the customer see when the AI hands the chat to a person?

The notice you chose in “Notice to the customer when handing off to a person”: the platform’s, your own text, a workflow’s or none.

Does the AI understand audio and photos?

Yes, if you turn on Voice note transcription and Image analysis in Tools. They use more credits than text.

Can I use the agent only at night?

Yes: in Service hours, **Custom** mode, with night windows (remember to split the night into two windows if it crosses midnight).

Can I test on WhatsApp without my customers seeing it?

Yes: turn on the **allowlist** with your number and write from your phone.

What happens if the AI credit runs out?

The AI pauses until the next cycle and leaves a note on each conversation so your team carries on by hand, unless you turn on overage with your Omnifox Balance.

Where do I see whether the agent answers well?

In **View quality** (with scoring on) and in **View analytics**.

Which engine level should I pick?

Start with Swift Fox. If it gets confused in complicated conversations, try Pro Fox. Max and Deluxe, for agents that handle delicate cases alone.

Are my documents deleted if I delete an agent?

No. Documents belong to the workspace and the other agents keep using them.

Need help?

Write to us from the **Support** button in the app or visit the help center at omnifox.io/help. If you are not using Omnifox yet, try it free for 7 days at **omnifox.io**.