AI Team — User Handbook

Control Center · Product version 2.0.0

Generated: 2026-07-10

00 Introduction

Welcome to the handbook for **AI Team** and its control environment, **Control Center**. On the following pages you'll learn how to manage your AI team, hand it work, review the results, and keep an eye on costs — no technical knowledge required.

What AI Team and Control Center are

You have a team of five AI specialists at your disposal, available around the clock. This is not a chatbot that forgets everything the moment you close the window — it's a team that works continuously: it processes tasks, watches over scheduled jobs (such as the morning briefing), remembers your preferences, and learns from your corrections.

Control Center is the web application where you manage this team. Think of it as the noticeboard of your AI company: it's where you assign tasks, chat with agents, review outputs, manage documents and the company wiki, track costs, and set up integrations (Telegram, CRM, WordPress, and more).

> 💡 **Tip:** You don't need to sit at your computer for the team to work. Tasks keep running after you sign out, and you can review the results whenever it suits you.

Your AI team in a nutshell

Role	Area	What they do for you
CEO	Director / coordinator	Your main contact — takes briefs, delegates to the team, reports results
CMO	Marketing	Copy, social media, email campaigns, brand, content
CSO	Sales	Outreach emails, follow-ups, proposals, sales materials
COO	Operations	Reports, processes, metrics, documentation
CTO	Technology	Automation, integrations, technical solutions

You'll find a detailed introduction to each role in the chapter Your AI Team.

> ✅ **Best practice:** The CEO is your main contact — they decide who to hand the work to and bring the results back to you. That said, you can also talk to the other agents directly.

Who this handbook is for

Business owners and employees who use the AI team in their everyday work. We assume no technical knowledge. Everything the product can do, you set up yourself — where the **Administrator** role is needed (e.g. integrations or user management), the text says so.

How to read this handbook

You don't have to read it from cover to cover. Jump straight to wherever the shoe pinches:

Situation	**Where to go**
I'm signing in for the first time, setting up my account	Getting Started
I don't know what each part of the app does	Finding Your Way Around
I want to talk to agents and assign work	Chat with Agents and Tasks
I want a recurring job (e.g. a weekly report)	Task Scheduler
I'm uploading source materials, looking for outputs	Data Center and Team Wiki
I want to message the team from my phone	Communication Channels
I want to know how much it costs	Costs and Limits
Something isn't working	Troubleshooting
I've run into an unfamiliar term	Glossary

Product version

This handbook describes **Control Center version 2.0.0** — the biggest update since launch (a new look with light and dark modes, live chat with multiple threads, a redesigned company wiki, voice control, and more). To find out which version you have, check the number at the bottom left of the menu; clicking it opens the **What's new** overview.

> ⚠️ **Watch out:** If your screen looks different from the descriptions in this handbook, check the version first. Older versions may have a different menu layout.

01 Getting Started

Your first sign-in, the setup wizard, and basic account adjustments — language, password, appearance. By the end of this chapter your environment will be ready and you'll know what to do on day one.

How to sign in

1. In your browser, open your Control Center address — you received it during handover (e.g. `https://your-company.team.cliqsales.com`).
2. In the **Username** field, enter the sign-in name you were given (usually your email address).
3. Enter your **Password** and click **Sign in**.

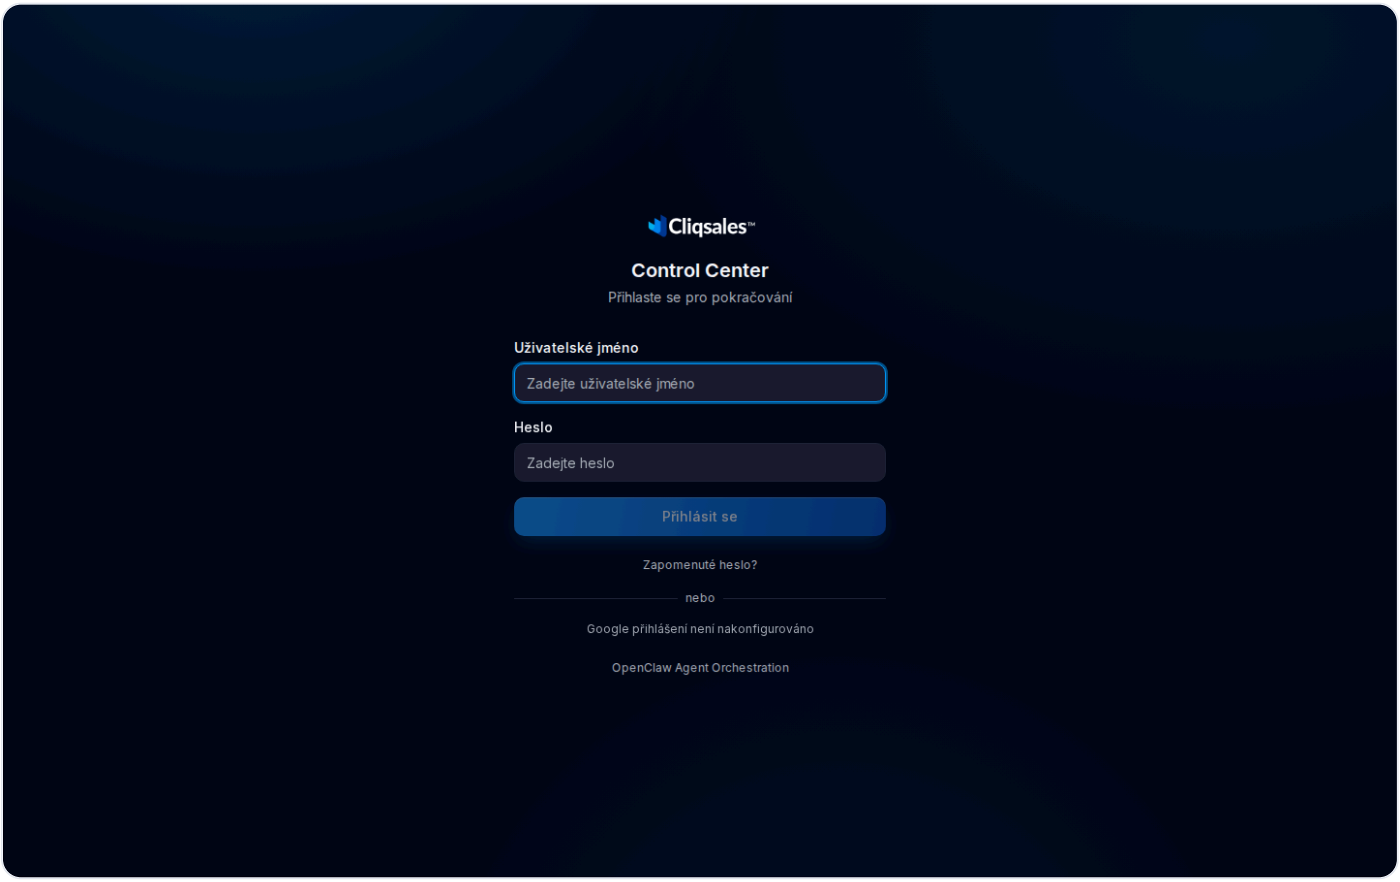

> ✅ **Best practice:** Bookmark your Control Center address in your browser. Every client has their own address — it's not one shared website.

Forgotten password

1. On the sign-in page, click the **Forgot password?** link.
2. On the **Password reset** page, fill in the **Username or email** field and click **Send reset link**.
3. You'll receive an email with a link to set a new password. **The link is valid for 2 hours** and can only be used once.
4. After clicking the link, enter your **New password** (at least 8 characters), repeat it in the **Repeat password** field, and confirm with the **Set new password** button.
5. The app redirects you to the sign-in page — sign in with your new password.

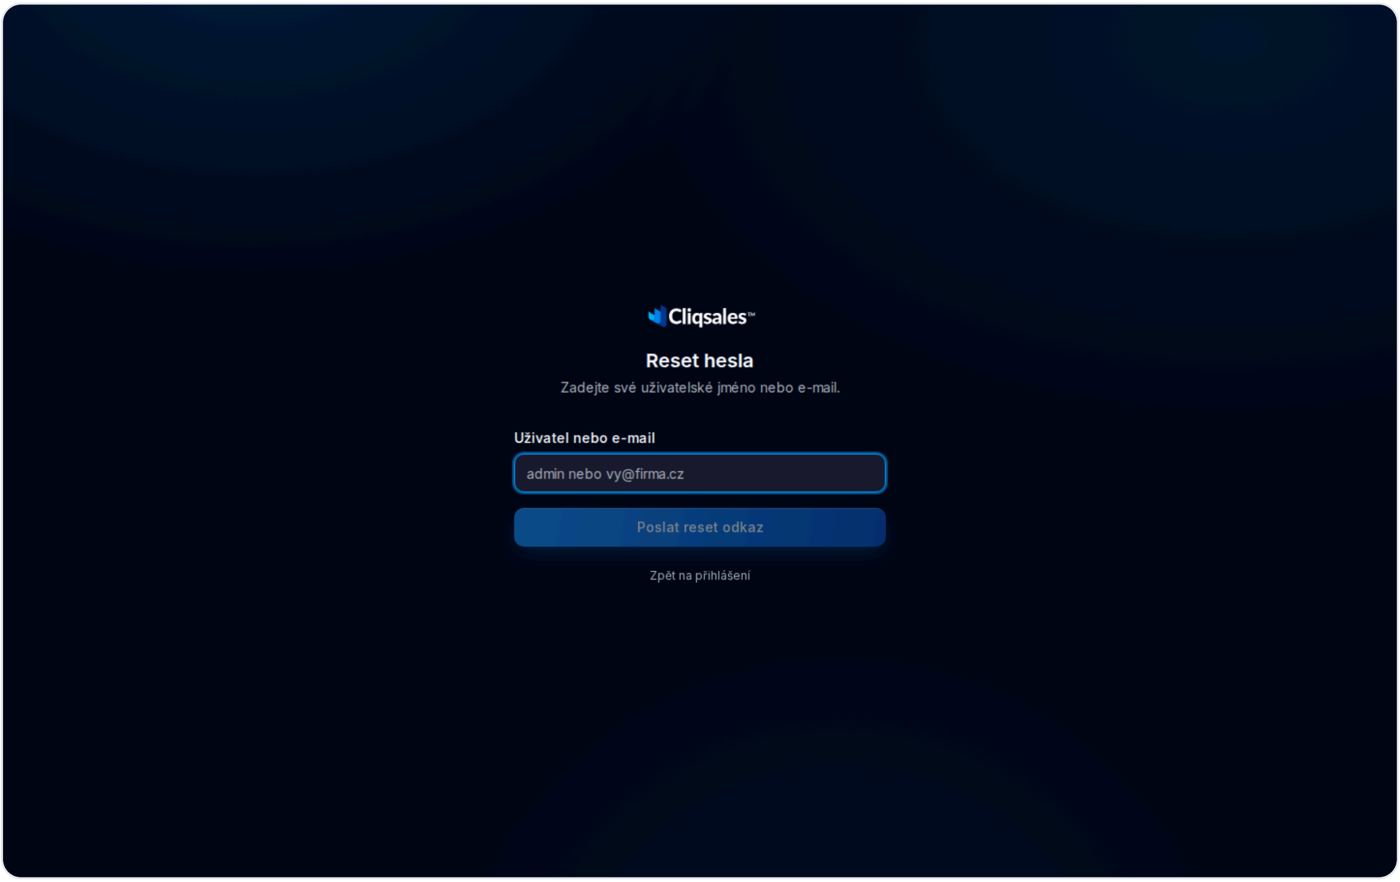

If something goes wrong:

- **The email didn't arrive** — check your spam folder and look for typos in the address you entered. For security reasons, the app always shows the same confirmation even if no account exists for the given address. If the email still doesn't arrive after a while, a colleague with the Administrator role can set a new password for you in **Settings → Users** (see Troubleshooting).
- **The link expired or has already been used** — the page tells you directly and offers a **Request a new link** button. Simply repeat the whole process.

First sign-in — the setup wizard

When you sign in for the first time, a **Welcome!** screen and the setup wizard open. In six steps it walks you through everything your AI team needs to know before it can start working. Your progress is saved — you can stop at any time and come back later.

Step	What you set up
1. **Set your own password**	Replace the handed-over password with your own — **at least 12 characters in this step**
2. **About you**	Name, how to address you, time zone, language, notes
3. **About the company**	Company, brand, tagline, website, industry, company description
4. **Operating context**	Budget, KPIs, key contacts
5. **Choose your AI provider**	The language model provider — OpenRouter is recommended; OpenAI, Anthropic, Moonshot/Kimi, and Minimax are also available
6. **Connect provider**	Paste and verify the API key (access key) of the selected provider

At the end you'll see a **You're all set!** screen with the quick actions **Chat with CEO**, **Create a task**, and **Go to dashboard**. Everything you filled in is handed to your AI team as its initial company context.

You don't have to finish everything right away:

- You can postpone an individual step with the **Skip this step** button.
- The **I'll finish later** option closes the wizard — it will remind you the next day.
- The × closes the wizard for good; you can relaunch it in **Settings** (the **Setup Wizard** section → **Restart onboarding**).
- Until you finish the wizard, an orange **Finish setup** button stays lit in the top bar — clicking it takes you back to where you left off.

> ✅ **Best practice:** Fill in the **About you** and **About the company** steps as honestly as you can. Everything you write here is what the agents use as their baseline knowledge of your business — it saves you from repeating yourself in chat.
>
> ⚠️ **Watch out:** If you don't have an API key yet, skip the step — you can add it any time later in **Settings → Integrations** (the step-by-step procedure is in the Integrations chapter). Without a connected model provider, however, the agents cannot work.

How to change your password later

1. Click the circle with your initials at the top right (**Account**).
2. Choose **Change password**.
3. In the **Change Password** window, fill in **Current password**, **New password** (at least 8 characters), and **Confirm new password**, then click **Save**.

In the same menu you'll also find **Sign out**.

How to switch the interface language

Control Center works in Czech, English, and Spanish.

1. Click the gear at the bottom left (**Settings**).
2. Open the **General** tab.
3. Under **Language**, choose **Čeština**, **English**, or **Español** and save your changes.

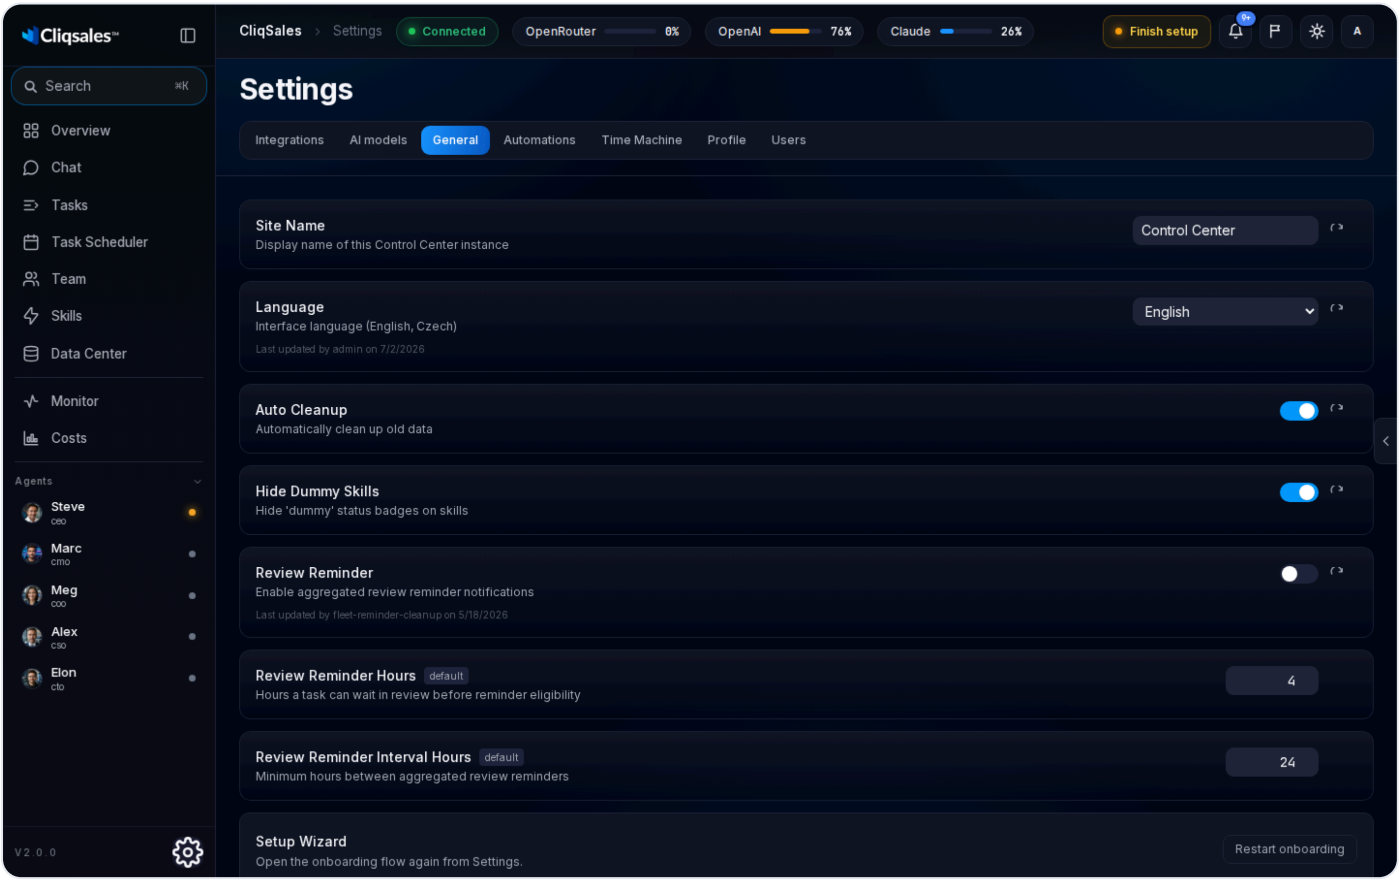

💡 **Tip:** The interface language is independent of the language you use with the agents. You can write to them in Czech even in the English interface — they'll understand and reply in Czech.

Light and dark mode

Switch the appearance with the **Switch color theme** icon (sun/moon) at the top right of the top bar. The choice only affects your browser — other users' appearance won't change.

Day-one checklist

The recommended start, so the team can work at full speed right away:

1. **Finish the setup wizard** — especially the steps about you and your company.
2. **Upload your core documents** to **Data Center → Documents**: a description of your products or services, price list, brand guidelines, templates. The steps are in the Data Center chapter.
3. **Introduce yourself to the CEO** in the **Chat** section: write who you are, what your company does, and what you expect from the team first. See Chat with

Agents.

4. **Assign a first small task** — e.g. "Prepare a short summary of our offering for the website". A small brief lets you get a feel for how the team works.
5. **Review the result and give feedback** — the team learns from your corrections and will hit the mark better next time.

> 💡 **Tip:** Connect Telegram — you can then task the agents from your phone, voice messages included. You can set up both the bot and the group yourself following Communication channels.

02 Finding Your Way Around

A quick tour of the screen: what's in the left menu, what the top bar shows, and what the right-hand live activity panel is for. After this chapter you'll know where to click, whatever you need.

The screen has three parts: the **menu** on the left (main navigation), the **bar** at the top (system status and your account), and an optional **Live Feed** panel on the right (what the team is doing in real time). After signing in you start on the **Overview** — a dashboard with the system status and team activity.

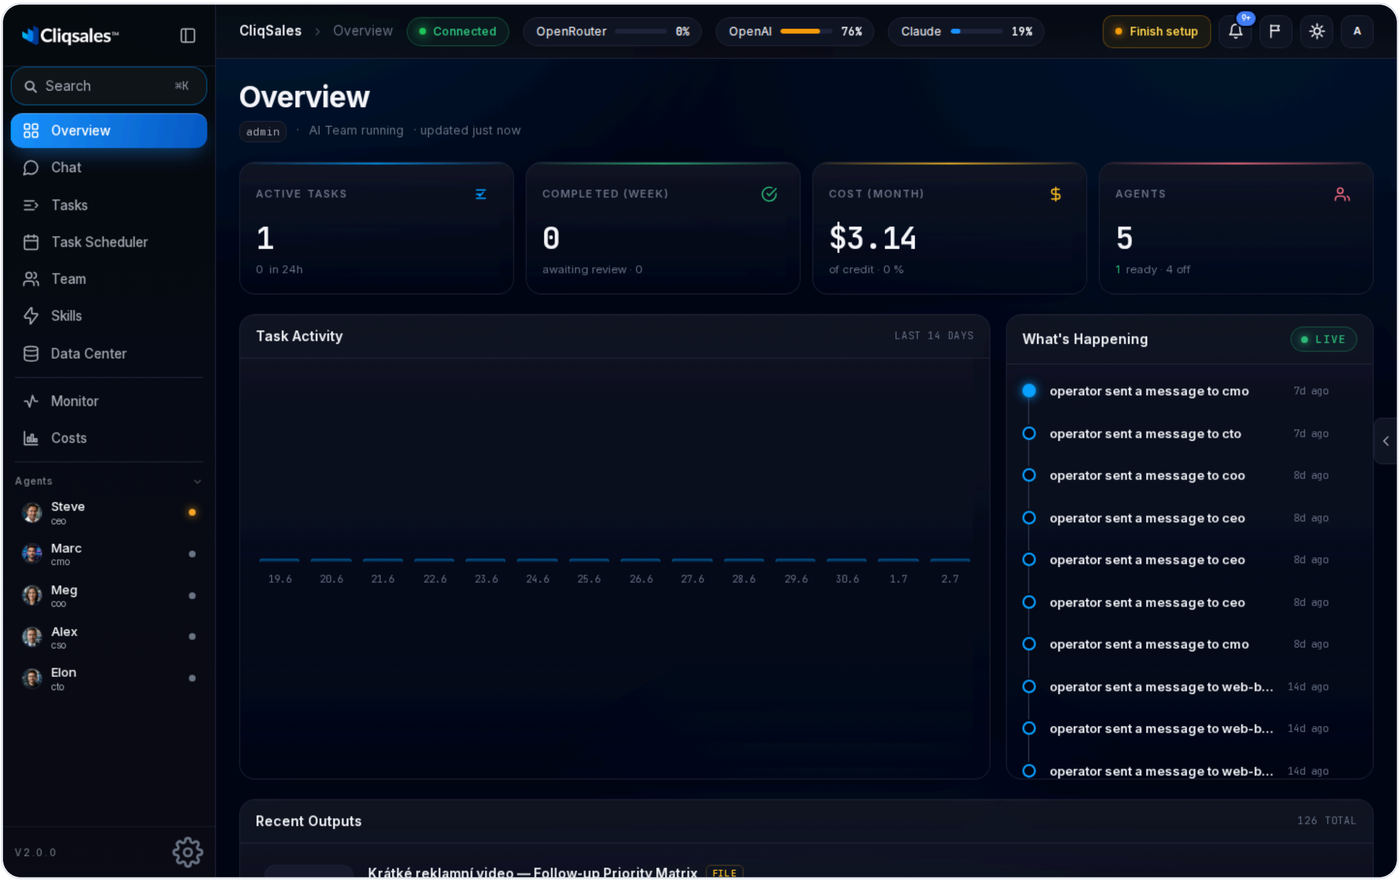

The left menu

Item	What it's for
Overview	The dashboard — system status and team activity in one place
Chat	Live communication with agents, conversation threads — see Chat with Agents
Tasks	The task board — see Tasks
Task Scheduler	Recurring and scheduled jobs (e.g. the morning briefing) — see Task Scheduler
Team	Your agents — status, detail, skills, history — see Your AI Team
Skills	The library of agent skills — see Skills
Data Center	Documents, outputs, and the team wiki — see Data Center
Monitor	Activity, logs, and sessions — mostly for advanced users and administrators
Costs	Subscription usage and pay-as-you-go costs — see Costs and Limits

Below the menu items is the **Agents** section — a quick list of your agents. Clicking an agent opens a chat with them directly. The colored dot next to the name shows their status: green = working right now, orange = waiting, red = error, grey = offline.

At the very bottom of the menu you'll find:

- the **version number** (left) — clicking it opens the **What's new** window,
- the **gear** (right) — opens **Settings**: company profile, language, Integrations, Automations, Time Machine (backups), users, and more.

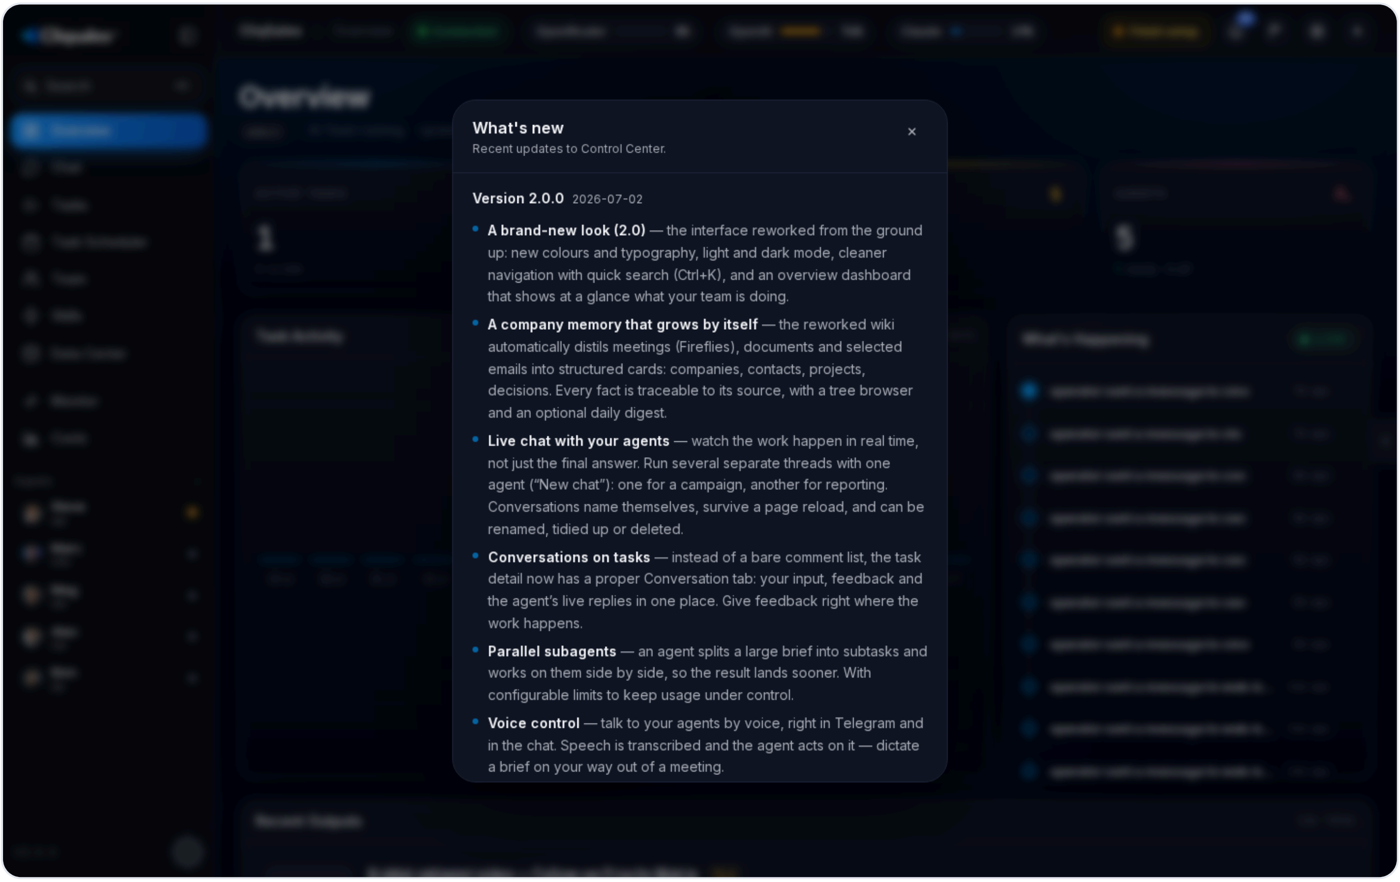

> 💡 **Tip:** The **What's new** window also lists older updates — click **Show previous versions**.
>
> ⚠️ **Watch out:** When the connection to the server is interrupted, the **Team** item may temporarily be greyed out with a **Requires gateway** note (the gateway is the connection to the AI core). It works again once the connection is restored — see Troubleshooting.

Collapsing the menu

You can narrow the menu to a slim strip of icons with the **Collapse sidebar** button next to the logo, or with the **[** key (it works whenever you're not typing in a text field). Expand the collapsed menu again by clicking the logo or pressing **[** once more.

Search (Ctrl/Cmd + K)

At the top of the menu is the **Search** field. Open it by clicking, or from anywhere with the **Ctrl + K** shortcut (**Cmd + K** on a Mac). Type at least 2 characters — it searches tasks, agents, conversation messages, team activity, and automation records. Clicking a result jumps straight to the relevant section.

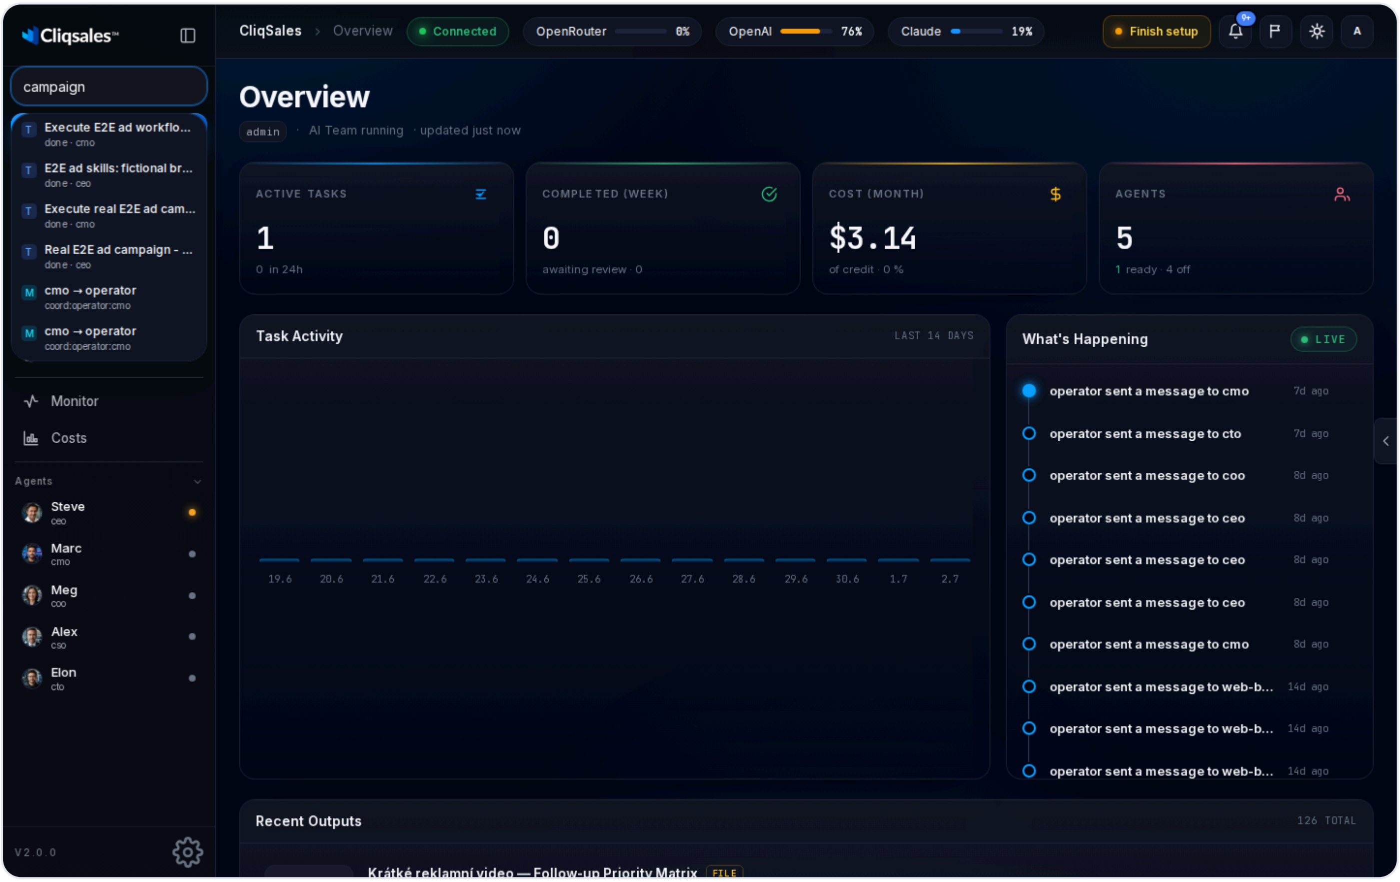

> 💡 **Tip:** Search is the fastest route back to work in progress: type part of a task name and you're in it — no clicking through the board.

The top bar

From left to right you'll find:

- **Breadcrumb navigation** — the app name and the section you're currently in.
- **Connection status** — a colored pill: green **Connected** means everything is running. Orange states (**Starting…**, **Connecting…**, **Restarting…**, **Updating…**) are transitional — wait a moment. Red **Disconnected** means the connection is down; see Troubleshooting.
- **Usage indicators** (e.g. OpenRouter, Claude, OpenAI) — how much of your subscription capacity or credit has been used. Details in Costs and Limits.
- The orange **Finish setup** button — only until you've finished the setup wizard (see Getting Started).
- The **bell** with the number of unread **Notifications** — opens the notification panel, where you can go through them and clear them with the **Mark All Read** button.

- **Feedback** — a window for quickly reporting a bug or an idea.
- **Switch color theme** — light/dark appearance.
- **Account** (circle with your initials) — **Change password** and **Sign out**.

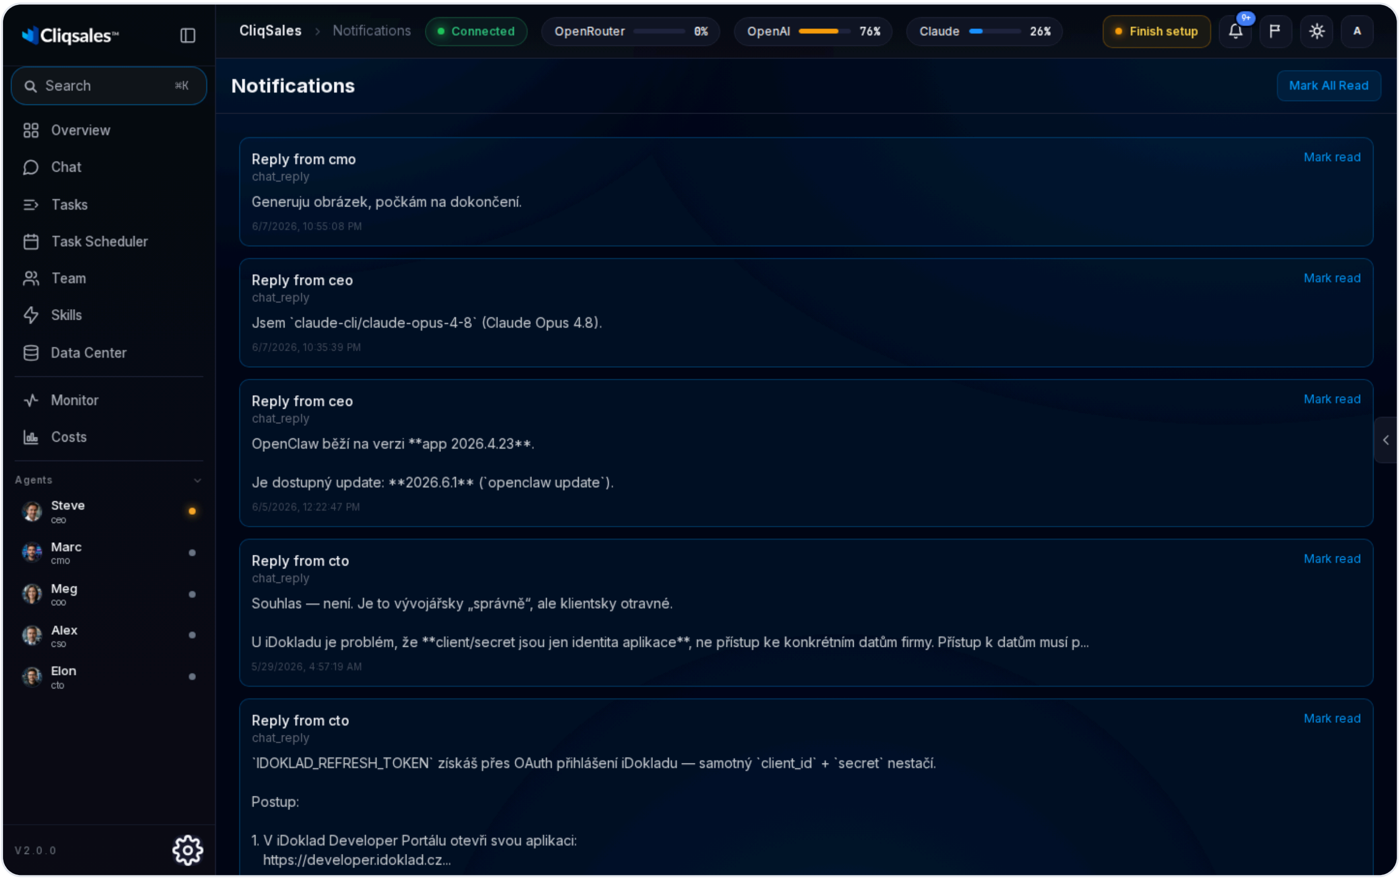

> ⚠️ **Watch out:** Clicking the green connection status pill restarts the AI core (hovering over it shows **Restart OpenClaw**). Only use it as directed by the Troubleshooting chapter, when the team has been unresponsive for a long time — the restart interrupts the agents' work for a few dozen seconds.

The right panel — Live Feed

On wider screens there's a **Live Feed** panel on the right: team events in real time (what the agents are doing, what has been completed) and an overview of active sessions (the agents' running work blocks).

- The arrow in the header (**Collapse feed**) narrows the panel to a slim strip.
- The × (**Close feed**) closes it completely.
- Reopen a closed panel with the small arrow button on the right edge of the screen (**Show live feed**).

> 💡 **Tip:** The Live Feed is a pleasant "peace of mind": you can see the team really is working, without having to open individual tasks.

Mobile view

On a phone, the left menu turns into a bottom bar with the most important sections: **Overview**, **Chat**, **Tasks**, **Team**, and **Data Center**. The other items are under the **More** button, which slides out the full menu. Connection status is shown by the colored dot at the top — green means everything is running.

The orange banner — missing model access

If an orange banner appears below the top bar saying *"No model provider access configured. Agents cannot process messages. Connect a model provider in Settings → Integrations."*, it means no model path is active (neither a subscription nor an API key) — the agents cannot work at that moment.

The fix: in **Settings → Integrations** (Administrator role), connect a model provider (e.g. OpenRouter).g. OpenRouter). Details in the Integrations and Troubleshooting chapters.

03 Your AI Team

Your AI team is made up of five specialists who are available around the clock. Each has their own role and strengths — and together they work like a small company: the director takes briefs, delegates them to colleagues, and brings the results back to you. In this chapter you'll learn who's who, how the team collaborates, what it remembers, and how to work with it in the **Team** section — including all the tabs in the agent detail, which you can set up yourself.

Who's who on the team

Agent	Role	Strengths
CEO	Director / coordinator	Your main contact. Takes briefs, distributes work across the team, and reports results.
CMO	Marketing	Copy, social media, email campaigns, brand, content.
CSO	Sales	Outreach emails, follow-ups, proposals, sales materials.
COO	Operations	Reports, processes, metrics, documentation.
CTO	Technology	Automation, integrations, technical solutions, specifications.

> ✅ **Best practice:** The CEO is your main contact — they decide who to hand the work to and bring the finished result back to you. You don't need to figure out which specialist is "the right one". Talk to the other agents directly when you're fine-tuning a detail in their field (e.g. polishing the tone of a text with the CMO).

How the team collaborates

The team works on the principle of **delegation**. When the CEO receives a brief that belongs to a specialist, they pass it on — and you'll see a handoff note in the chat (e.g. "CEO handed off to CMO"). The CEO keeps you informed about progress and results; you don't have to forward anything yourself.

For larger briefs, an agent can **split the work into subtasks and run them in parallel** — spinning up temporary helpers (so-called subagents) that work at the same time. There's nothing to set up or switch on; the system keeps the number of concurrent helpers in check by itself. You just follow the progress in chat or in the task detail. You'll find the toggle and the run history in the agent detail on the **Subagents** tab (described below).

> 💡 **Tip:** You don't have to slice a big brief into pieces yourself. Write it as a whole ("Prepare the complete launch package for the product: website copy, 5 social posts, an outreach email") — the team will divide the work more efficiently than if you created five separate small tasks.

How the team remembers

Your AI team is not a chatbot that forgets everything once you close the window. Its memory does have several levels, though, and it pays to know where things live:

- **The agent's working memory** — each of the five agents keeps their own notes: priorities, constraints, and decisions to keep in mind. You'll find them in the agent detail on the **Memory** tab — you can read and edit them.
- **Ongoing consolidation** — automatic memory maintenance runs in the background: what an agent uses repeatedly while working is gradually added to their permanent notes. There's nothing for you to run.
- **Long-term company memory** — held by the **team wiki** (ch. 8) and your source materials in Data Center. This is where things that should survive for months are stored.
- **Context about you and your company** — the profile on the **User** tab (company, product, target audience, tone of voice). It syncs to all agents at once.

There is no automatic "telepathy" between agents. What you tell one agent in chat, the others don't know on their own. Information spreads through the team in three ways: via the team wiki, via documents in Data Center, and via task handoffs. If you want the whole team to know something permanently, it belongs on the **User** tab (company facts), in the wiki, or in the documents — not just in a chat with one agent.

Where things are stored

Type of information	Where it lives	How to check / edit it
Facts about you and your company (colors, tone, target audience)	the **User** profile — shared by all agents	agent detail → **User** tab, **Save & Sync to All Agents** button
One agent's priorities and decisions	the agent's working memory	agent detail → **Memory** tab (read and edit)
The agent's personality and working rules	the agent's **Soul**	agent detail → **Soul** tab
Company knowledge (projects, meetings, conclusions)	the **team wiki**	**Data Center** section → **Wiki** tab (ch. 8)
Your source materials and files	documents	**Data Center** section → **Documents** tab (ch. 7)
The team's finished outputs	outputs	**Data Center** → **Outputs** tab, or agent detail → **Outputs**
An ongoing conversation	the chat thread	the Chat section — applies only within that thread

> ✅ **Best practice:** Correct specifically. "Not this generic tone — write more personally and informally" teaches the agent more than "I don't like it". The more you work with the team, the better it knows you.
>
> 💡 **Tip:** Agents can also forget. Say "forget that I prefer blue" — the agent updates its notes, and you can check the result on the **Memory** tab of its detail.
>
> ⚠️ **Watch out:** The team only knows what it has been given — from chat, documents, the wiki, or the **User** profile. If an agent has forgotten something (it happens rarely), simply repeat the information; and if it should hold forever, store it where it belongs per the table above.

The Team section

You'll find the overview of the whole team in the left menu under **Team**. At the top you can see a summary (how many agents you have, how many are ready, and how many are off) and a switch between two views:

- **Grid** — agent cards: photo, name, role, assigned AI model, and quick **Task** (a new task for the agent) and **Chat** (write directly) buttons.
- **Org** — the organizational structure: who reports to whom.

Clicking an agent's card opens their detail.

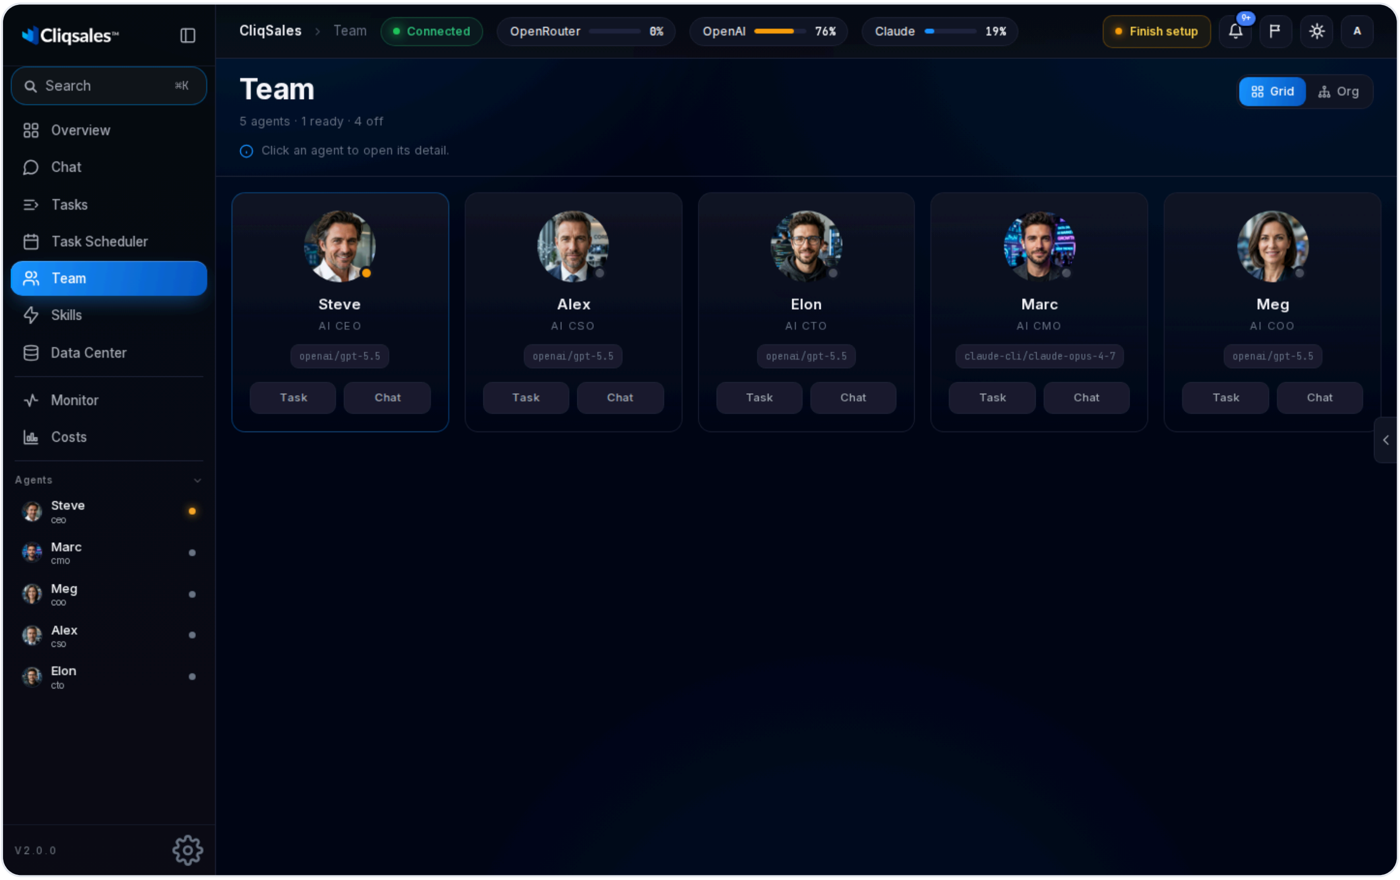

What the colored dots mean (agent status)

Next to each agent's photo (in the **Team** section and in the **Agents** list at the bottom of the menu) there's a colored dot:

Color	Status	Meaning
🟢 green	**Busy**	the agent is currently processing a task or message
🟡 yellow	**Idle**	the agent is ready and waiting for work
⚪ grey	**Offline**	the agent isn't running right now
🔴 red	**Error**	something went wrong — see Troubleshooting

How to wake an agent

Agents normally wake up on their own — when you write to them or when scheduled work is waiting for them. If an agent is **Offline** and you need them right away:

1. Open the **Team** section and click the agent.
2. In the detail header, click **Wake Agent**.
3. The agent wakes up and checks whether any work is waiting for them.

Agent detail

The agent detail is a team member's "personnel card". At the top you can see the photo, name, role, the AI model in use, who the agent reports to (**Reports to**), and the current status with a wake button; next to it, the **Reset Session** button reloads the agent's configuration and starts with a clean slate (useful after bigger edits — see Soul below). Go back with the arrow at the top left (**Back to squad**).

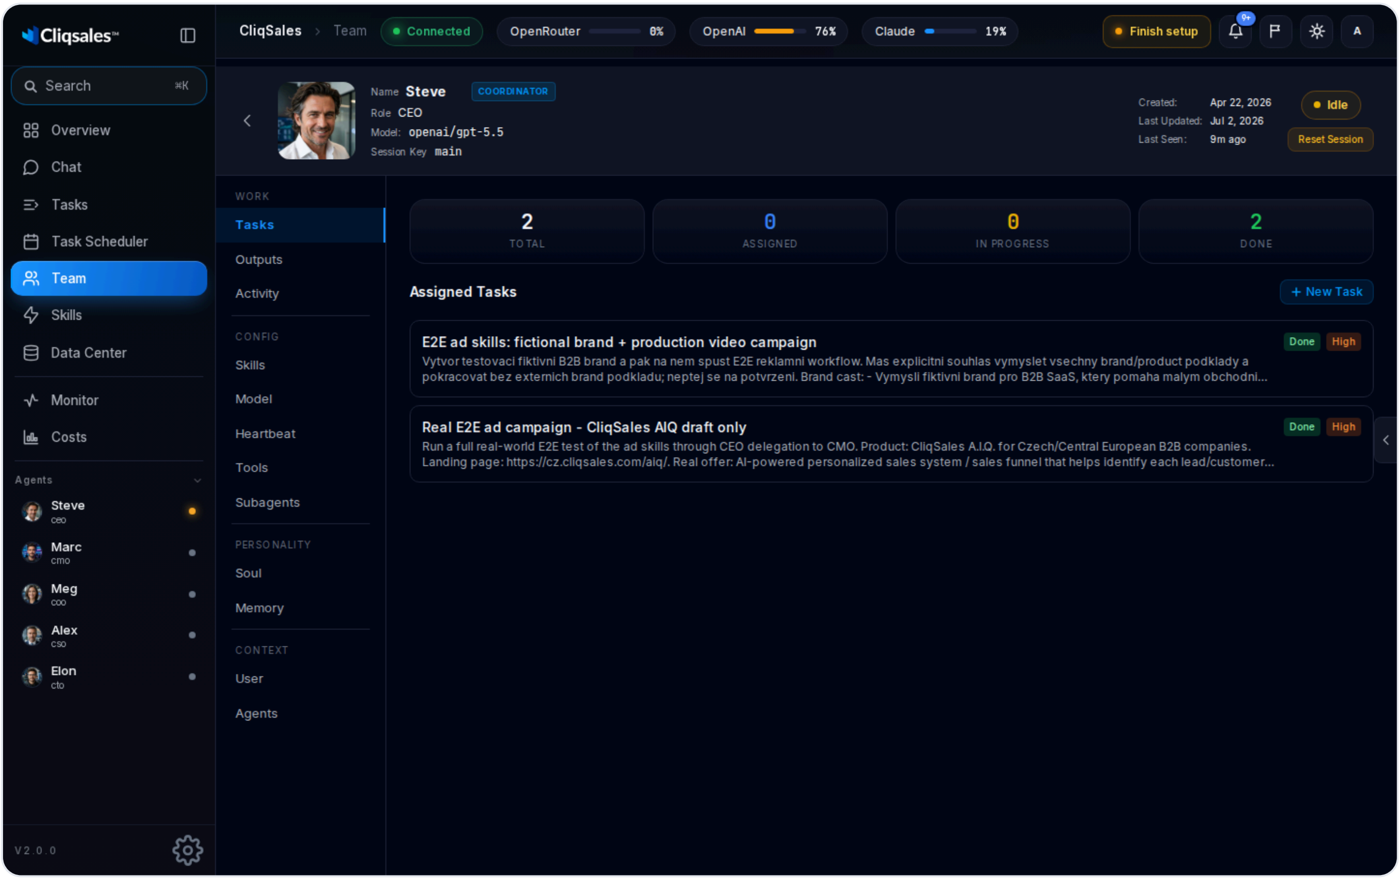

On the left are tabs organized into four groups. **WORK** covers everyday operation: **Tasks** (the agent's tasks; from here you create a **New Task** — details in the Tasks chapter), **Outputs** (everything the agent has created; view or download) and **Activity** (what the agent did and when). The remaining groups — **CONFIG**, **PERSONALITY**, and **CONTEXT** — are covered in the sections below; you can set up all of it yourself.

The **Viewer** role is enough for browsing; to make changes you need **Operator** or **Admin** (roles are explained in the Administration and security chapter).

Skills — what the agent can do

The agent's skills: what's enabled and what you can switch on or off (details in the Skills chapter). Changes are saved immediately but take effect the next time the agent wakes up — work in progress isn't interrupted.

Model — which AI model the agent uses

The **Model** tab (headed **Model Assignment**) determines which AI model the agent "thinks" on. You'll see a **Current:** row with the assigned model and a list of available models grouped by provider. You pick one primary model; **fallback models** for when a limit runs out are set globally in **Settings → AI models** (**Admin** role). The list tells you:

- **The price** — a figure like `$3/$15` is the price per 1 million tokens, input/output in USD. A cheaper model = cheaper operation.
- **Subscription limits** — for providers with time windows you can see when the limit renews (e.g. "hourly: resets in 32 min").
- **Unavailable models** — providers marked "not configured in Integrations" need to be connected first in **Settings → Integrations**; their models can't be selected. Hidden providers unfold with **Show more providers**.

To change the model: pick one from the list (you can filter with the **Search models…** field), and a **Selected:** bar appears — confirm with the **Save** button, or back out with **Discard**. The saved model is written straight into the agent's runtime and is used for their next work.

> 💡 **Tip:** Choose the model by role: an agent who writes important copy deserves a more expensive, smarter model; an agent doing routine overviews gets by with a cheaper one. You track costs by model in the Costs section.
>
> ⚠️ **Watch out:** A badly chosen model is the fastest way to an expensive or slow team. Before changing it, note down the name from the **Current:** row; after the change, verify the result on a small brief in chat — and if quality drops, go back to the original model.

Heartbeat — the agent's regular wake-up call

Besides your messages, agents also wake up on their own at regular intervals — that's called the heartbeat. The **Heartbeat** tab (headed **Rules**) contains a short checklist the agent reads and carries out on such a wake-up. By default it's practically empty: the agent just confirms it's alive and goes back to sleep. Editing it (the **Edit** and **Save** buttons) makes sense when you want a small recurring check — e.g. "on wake-up, check whether any unanswered task comments are waiting".

> 💡 **Tip:** For scheduled work at a specific time ("every Monday at 8:00, prepare the report"), the Task Scheduler is a better fit — leave the heartbeat for small ongoing checks.
>
> ⚠️ **Watch out:** Every wake-up costs tokens. Don't put big tasks or long lists into the heartbeat — you'll increase consumption without getting better results (see Costs).

Tools — an overview of the agent's environment

The **Tools** tab (headed **Environment**) is a notes sheet about the tools and services available to the agent — where to find their skills, where they store files, and so on. The agent reads it as orientation notes. Important: this text does **not control** what the agent is actually allowed to do — real capabilities are switched on and off on the Skills tab. So treat Tools as an overview; you normally don't need to edit it.

Subagents — parallel helpers

The **Subagents** tab controls whether the agent may spin up temporary parallel helpers for big briefs (see "How the team collaborates" above). It switches with a single **On** / **Off** toggle. Below the toggle you can see the helpers' recent runs and their status (running / done / failed); until the agent has started any, it reads "No subagent runs yet."

Two things are not set here:

- **The helpers' limits and default model** are shared by the whole team — in **Settings → AI models** (**Admin** role) you'll find **Max total subagents**, **Max subagents per agent**, **Run timeout (s)**, and **Default subagent model**. Also, a helper never starts a helper of its own — the work doesn't branch any deeper.
- **Agents on Claude models** use their own built-in helpers automatically — for them the tab shows a **Native** badge instead of the toggle.

> ⚠️ **Watch out:** More concurrent helpers mean a faster result, but also more tokens consumed at once. If costs are a concern, leave the limits at their defaults.

Soul — personality and working instructions

Soul is the agent's "soul": a text (headed **SOUL Configuration**) that defines who the agent is and how they should work. It typically covers the agent's role and mission, tone and communication style, when to act, what they do on their own, escalation rules ("when to ask a human"), and forbidden things ("never send anything out without approval"). The text has a limit of 20,000 characters — the counter sits right by the heading. While editing, there's also a **Select a template...** picker at the top with a **Load** button that loads a ready-made personality template into the editor.

Editing is safe when you know what you're touching:

- **Feel free to edit:** tone and language ("write more concisely", "communicate in English"), output format ("reports always as bullet points with numbers"), priorities ("quality over speed"), and preferred ways of working.
- **Change with care:** the agent's role, delegation rules (who hands what to whom), and safety principles (approvals, prohibitions). This is where you can accidentally "rewrite the agent's character".

The proven approach: **small change → test → revert if needed.**

1. On the **Soul** tab, click **Edit** and first copy the original text somewhere safe (as a backup).
2. Make one small change and click **Save**.
3. Test the agent in chat on a typical brief.
4. Not right? Paste the backup back and save. Working? Continue with the next edit.

> 💡 **Tip:** Before you touch the Soul, try telling the agent the instruction in chat, or writing it into the **User** tab (company facts) or **Memory** (the agent's notes). Edit the Soul only when the change should hold permanently and for all of the agent's work.
>
> ⚠️ **Watch out:** A saved change takes effect the next time the agent wakes up — you can force it immediately with the **Reset Session** button in the header (the agent then starts with a fresh conversation). And the template's **Load** button overwrites the entire draft text — hence the backup of the original wording before you start cutting.

Memory — the agent's working memory

The **Memory** tab is the agent's notebook: longer-lasting priorities, constraints, and decisions to remember across conversations. The agent fills it in on their own while working (and the ongoing consolidation in the background keeps it maintained), but you can see into it and edit it — use the **Edit** button to fix an outdated note or add a new one.

> ✅ **Best practice:** When you tell an agent "forget that..." or correct them, take a look here — you'll see in black and white what they actually wrote down.

User — what the team knows about you and your company

The **User** tab shows the shared client profile: structured fields in the groups **About you** (name, how to address you, language), **About the company** (company, website, description), **Offer and audience** (products, target audience, communication style, brand colors), and **Operating context** (budget, KPIs, contacts).

The profile is **shared by the whole team** — it's the same profile you'll find in **Settings**, just at hand right next to the agent. After editing, click **Save & Sync to All Agents**: the context for all agents is regenerated from the fields at once (confirmed by the "Synced to N agents" message).

> 💡 **Tip:** This is the right place for facts the whole team should know — brand colors, tone of voice, target audience. Writing them here is more reliable than telling each agent separately in chat.

Agents — how the team is put together

The **Agents** tab is an informational overview of the internal organization: the **Team Manual** describes who the agent works with and how the team is organized (read-only), and below it is the **Direct Reports** list — for the CEO you'll see the whole team here, for specialists usually "No direct reports". Clicking a report takes you to their detail. There's nothing to set here; it's for quickly seeing who reports to whom (the **Org** view in the **Team** section shows the same thing graphically).

How to upload an agent's photo (avatar)

You can give the agents faces of their own — the photo then appears in chat, in the **Team** section, and in the menu.

1. Open the agent detail and click their photo at the top left.
2. The **Update Avatar** window opens.
3. Drag an image into the marked area, or click **Choose File**.
4. Supported formats: PNG, JPG, WebP, or GIF.

> 💡 **Tip:** Custom team avatars help you orient yourself quickly — in chat you can see at a glance who is replying to you.

04 Chat with Agents

Chat is the fastest way to work with your AI team: you write a message as you would to a colleague and watch in real time what the agent is doing — not just the final answer. In this chapter you'll learn to use the chat (threads, attachments, voice messages, search) and, above all, how to phrase briefs so the results are worth it.

Where to find the chat

- In the left menu, click **Chat** — a full screen opens with the list of agents and conversations on the left.
- Or click a specific agent in the **Agents** section at the bottom of the left menu — the conversation with them opens directly.
- A third way: in the **Team** section, every agent card has a **Chat** button.

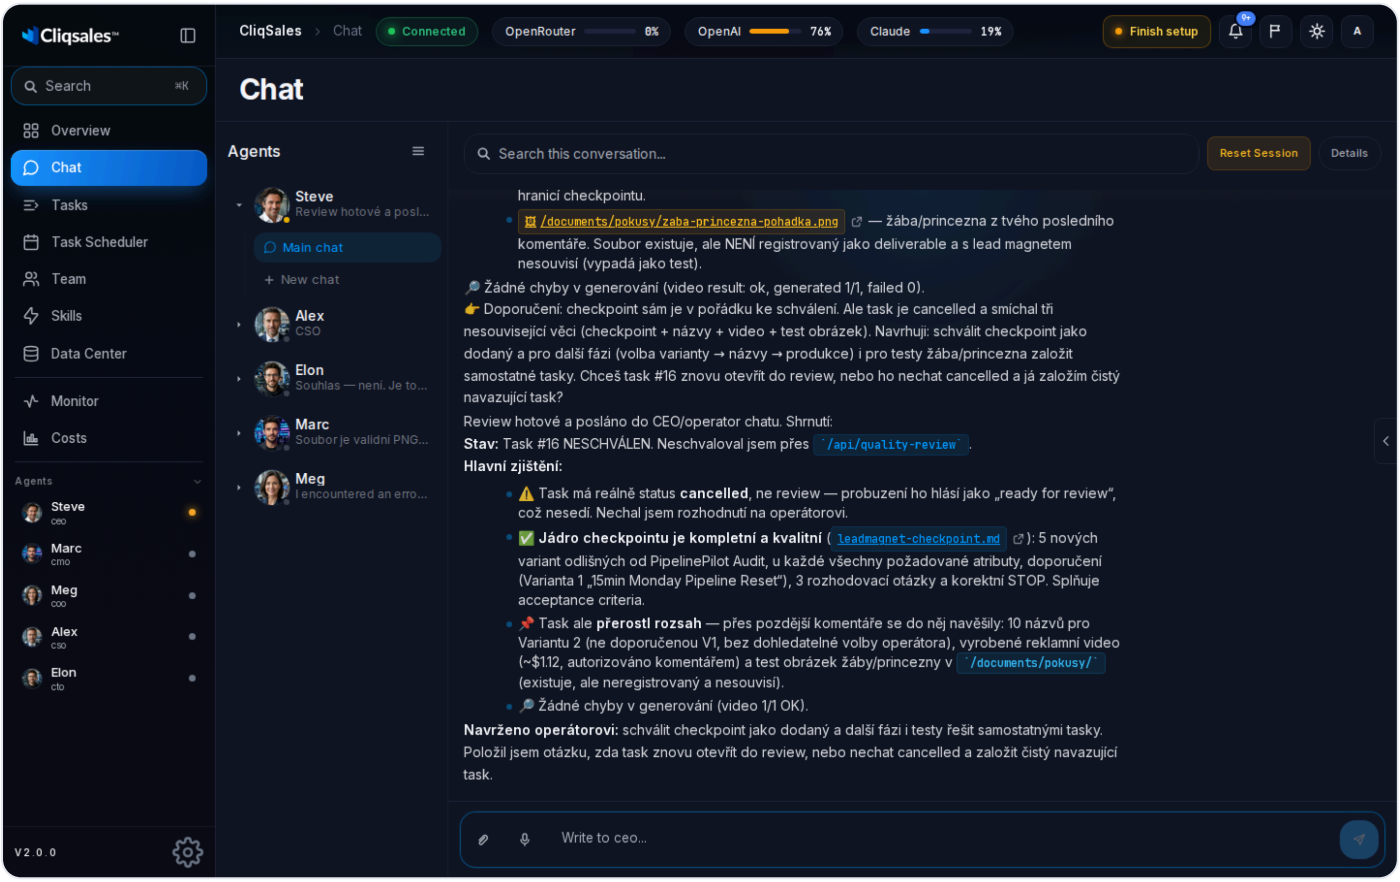

Send a message with **Enter**; insert a new line with **Shift + Enter**.

Live work progress

When an agent works on something bigger, you don't just see "typing…" but ongoing **step bubbles**: "Reading context…", "Saving output…", "Generating an image…" and so on. You can see that something is really happening — and what exactly.

The steps of a single run are stacked into one group:

- The group header shows the status — **In progress**, **Completed**, **Needs attention**, or **Failed** — the number of steps (**Steps: N**), and a stopwatch while work is running.
- A collapsed group shows only the last step; the **[+]** button expands all steps, **[−]** collapses them again.
- Each step has a **Details** / **Hide details** link with technical specifics (result, duration, any error).

> 💡 **Tip:** The **Details** button above the conversation turns on details for the whole conversation at once. It's handy when you're tracing exactly what happened during a run — otherwise leave it off; the conversation is cleaner.
>
> ⚠️ **Watch out:** The **Needs attention** status doesn't always mean a problem — the agent often hits a warning and continues via a fallback path. Watch whether the work ended in **Completed** and what the final reply says.

Threads: multiple conversations with one agent

You can hold several separate conversations (threads) with each agent — say, one thread for a campaign and another for reporting. Threads survive a page refresh and a restart.

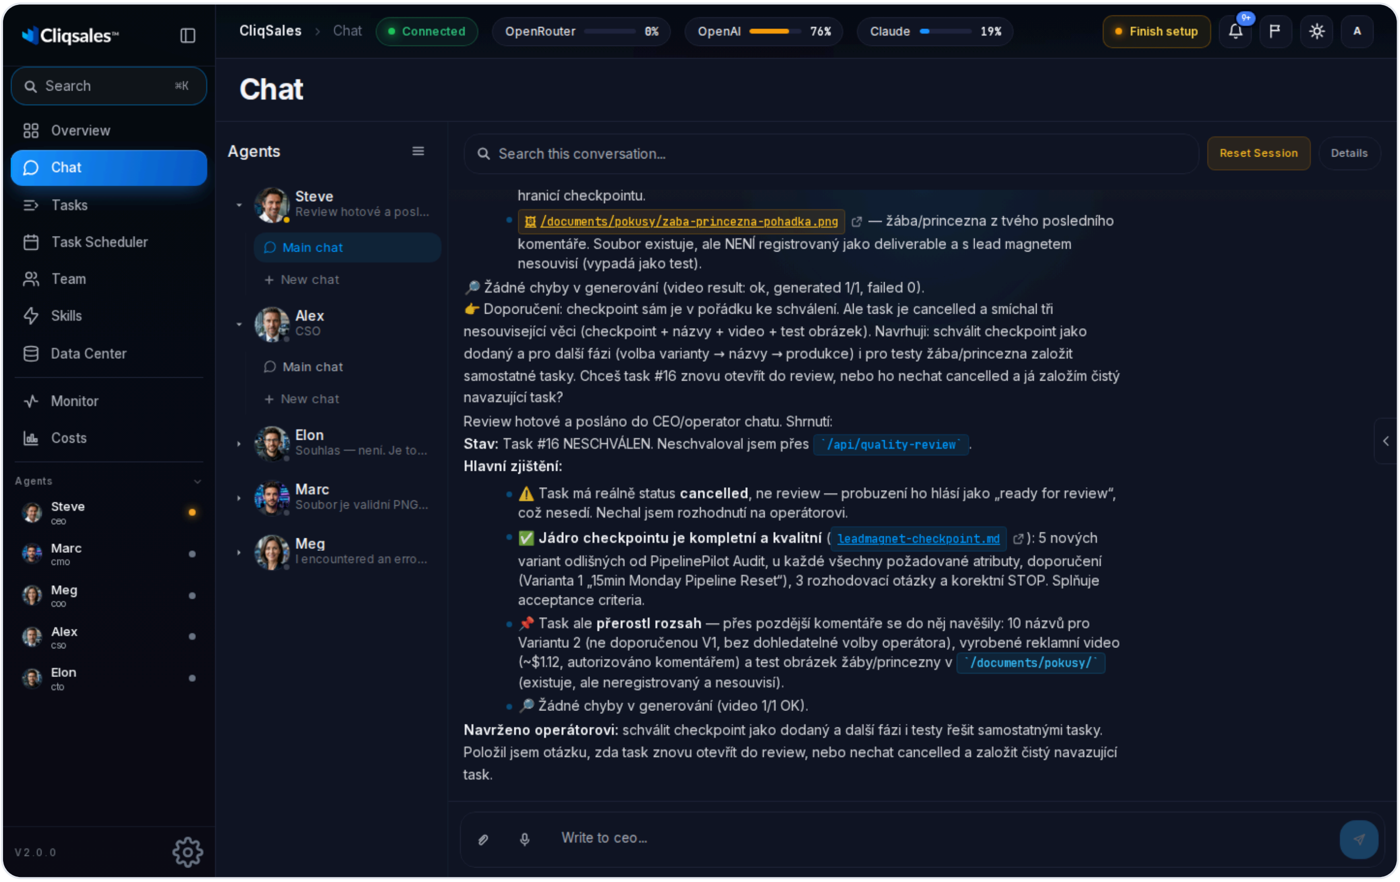

How to start a new thread

1. In the list on the left, click the agent's row — the arrow expands their threads.
2. Click **New chat**.
3. The new thread opens as **Untitled** and names itself based on the content after your first message.

Every agent always has a default **Main chat** thread — it can't be deleted or renamed. Chats for running tasks also appear among the threads (rows starting with `#number`); these are managed by the system and disappear from the list once the task is approved or cancelled.

Renaming, resetting, and deleting a thread

1. Hover over a thread and click the ··· button (**Thread options**); you can also open the menu with a right-click.
2. Choose:

- **Rename** — enter the new name and confirm with Enter.
- **Reset** — clears the agent's working context for this thread; the messages stay, but the agent starts "with a clean head".

- **Delete** — removes the thread entirely.

> ⚠️ **Watch out:** Use **Reset** when a conversation has become "tangled" and the agent keeps going in circles. Keep in mind, though, that the agent forgets the context of that thread — repeat any important instructions afterwards.
>
> ✅ **Best practice:** One thread = one topic. A campaign, reporting, and website ideas belong in three threads, not one. The agent keeps the context straight and you'll find your way around older conversations easily.

Attachments: send files to an agent

You can attach **up to 4 files** to a single message:

Type	Limit	Formats
Images	up to 5 MB	PNG, JPG, WebP, GIF
Documents	up to 20 MB	PDF, DOCX, TXT, MD
Audio	up to 100 MB	MP3, WAV, M4A, AAC, OGG, FLAC
Video	up to 100 MB	MP4, WebM, MOV, M4V

How to add an attachment:

- Click the paperclip (**Attach file**) next to the message field and pick the files, **or**
- simply **drag the files** into the lower part of the chat — a "Drop files to attach" frame appears, **or**
- paste an image from the clipboard with **Ctrl + V** straight into the message field.

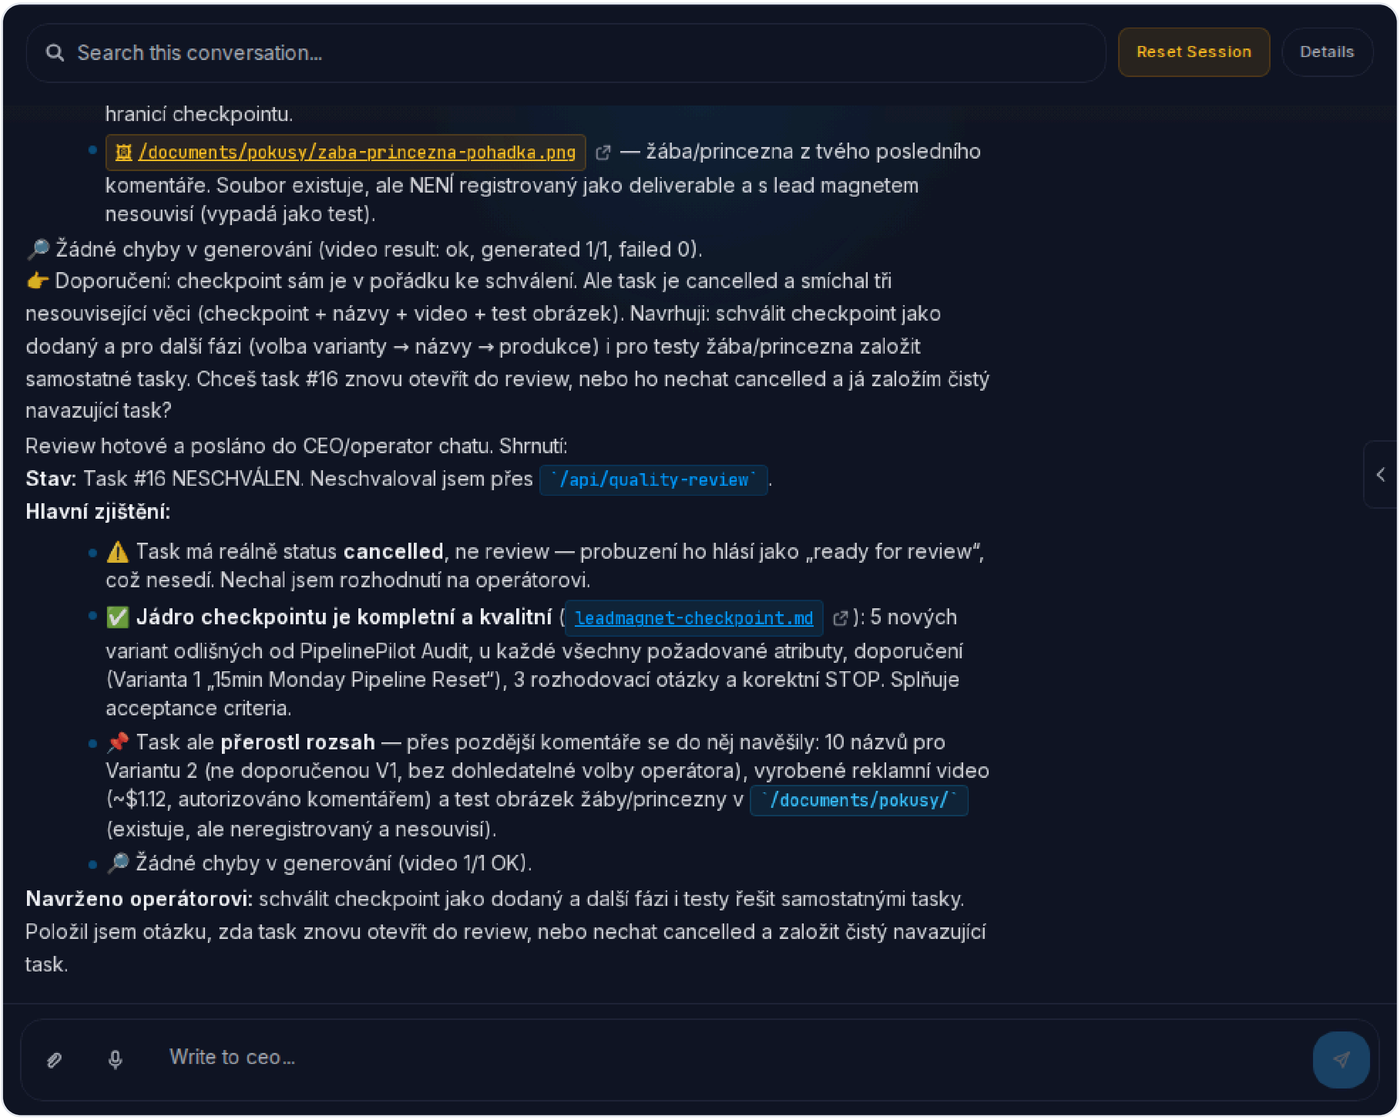

> 💡 **Tip:** Materials you'll use repeatedly (price list, brand manual, product description) are better uploaded to Data Center — there the whole team has them permanently available and you don't have to attach them again.

Voice messages

You can speak a brief instead of typing — the system transcribes it to text and the agent acts on it:

1. Click the microphone icon (**Record a voice message**).
2. Speak; a timer runs alongside a **Recording** indicator.
3. Clicking the send button (**Send voice message**) finishes the recording and sends it straight away. The × (**Cancel recording**) discards it.

> ⚠️ **Watch out:** The first time you use it, the browser asks for microphone access. If you see the message "Microphone unavailable or access denied.", allow the microphone for your Control Center address in your browser settings.
>
> 💡 **Tip:** You can also give the team voice briefs from your phone via Telegram — see Communication Channels.

Searching within messages

Above the conversation is a field with a magnifier (**Search this conversation...**):

1. Type the search term — matches are highlighted and a counter shows "1 / N".
2. Use the arrows to jump to the **previous** / **next match**.
3. **Clear search** removes the highlighting.

> 💡 **Tip:** Search goes through the messages loaded in the window. If you're looking for something older, first click **Load previous 100 messages** at the top of the conversation — and then search again.

File links in replies

When an agent mentions a file in a reply (an output, a document), it appears as a colored "chip" with the file name:

- Clicking the chip opens a preview right in the chat (images, PDFs, texts).
- The small icon next to the chip (**Open in Data Center**) opens the file in Data Center, where you can keep working with it.
- Images and videos the agent created are displayed directly in the chat as previews.

When an agent takes a long time

Simple answers take seconds, bigger tasks minutes. What to do when nothing seems to be happening:

1. Look at the progress bubbles — if steps keep appearing, the agent is working. On large briefs it splits the work into parallel subtasks, so it can handle more at once.
2. If a message appears saying the message is still being processed or that the agent didn't return a reply within the time limit, the work usually keeps running in

the background — the outputs will be filled in; refreshing the page also helps.
3. Check the agent's status (the colored dot) in the **Team** section or in the **Agents** list — see Your AI Team.
4. If nothing happens even after a longer while, follow the Troubleshooting chapter.

> 💡 **Tip:** You don't have to sit by the chat. The message field stays active, you can draft your next message — and come back to the result later; the progress and the reply stay in the conversation.

How to phrase a brief

The quality of the result stands and falls with the brief. A good brief has three parts: **context** (for whom, for what), **goal** (what exactly should be produced), and **format** (length, tone, structure).

```
❌ "I need a LinkedIn post"

✅ "I need a LinkedIn post about the launch of product XYZ.
    We're targeting HR managers at companies with 50–500 employees.
    Tone: professional but approachable. Max 150 words.
    End with a question to spark discussion."
```

> ✅ **Best practice:** Give the team source materials along with the brief — upload the relevant documents to Data Center in advance, mention the deadline, and write down what you know about the situation. The more context, the less guesswork.
>
> ✅ **Best practice:** Iterate. The first version doesn't have to be perfect — and that's fine. Send specific feedback ("Good, but cut it by a third and drop the jargon") and the agent reworks the result. Two quick iterations beat an hour spent polishing the "perfect" first brief.
>
> 💡 **Tip:** When you're not sure how to approach a brief, ask the CEO: "What do you need from me to prepare X?" The agent will ask for the missing information.

New thread, or keep going?

- **Continue in the thread** when you're following up: feedback on a draft, additions to a brief, the next step of the same project. The agent has the context.
- **Start a new thread** when you're opening a new topic or a new project. A long conversation about everything slows the agent down and dilutes the quality of the answers.

Chat, or a task?

- **Chat** is for quick things: questions, brainstorming, small texts, fine-tuning details, feedback.
- A **task** is for bounded work with outputs: it has a brief, an assigned agent, a status on the board, downloadable outputs, and an approval step — details in the Tasks chapter.

You don't have to decide up front, though: when you give the CEO a bigger brief in chat, they'll create the task and delegate it themselves.

Common pitfalls

> ⚠️ **Watch out — vague briefs:** "Do something about our marketing" ends in a generic result. Always state the goal, the audience, and the format; otherwise the agent guesses.
>
> ⚠️ **Watch out — mixing topics in one thread:** When one conversation covers a campaign, invoicing, and a website idea, the agent loses the thread and the answers get worse. Split topics into threads.
>
> ⚠️ **Watch out — expecting knowledge the team never got:** An agent doesn't know the content of a meeting you didn't tell it about, or a document you didn't upload. Before blaming the result, check that the team had the source material — in chat, in Data Center, or in the wiki.
>
> ⚠️ **Watch out — blind trust in facts:** AI can "make up" a fact (hallucinate). Always check the numbers, names, and claims in important outputs — the final decision is yours.

05 Tasks

The **Tasks** section is your AI team's work board. It shows everything the team is working on, what's waiting for your review, and what's already done. This is where you'll come back most often: it's where you assign tasks, track them, and approve their results.

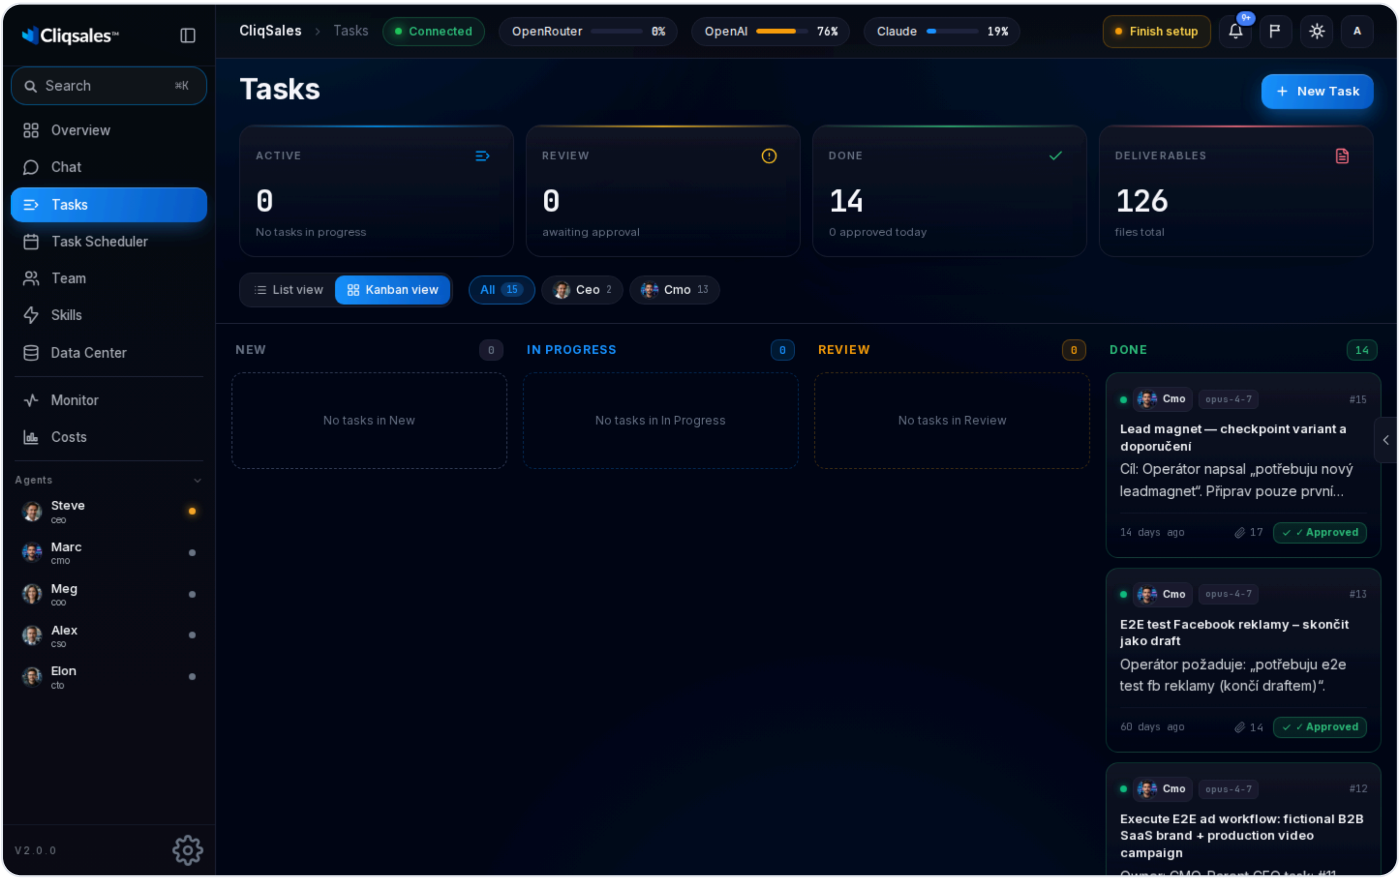

How to read the board

The board has 4 columns that a task moves through:

Column	Meaning
New	The task is waiting to get going — it was just created or assigned to an agent
In Progress	An agent is working on the task right now
Review	The result is waiting for a check — yours, or an internal quality check
Done	The task is finished and approved (tasks that failed also land here — you'll recognize them by the red dot and the **Failed** status)

Each card shows the assigned agent, the task number, the title, a snippet of the description, the number of outputs, and the time of the last change. Approved tasks carry a green ✓ **Approved** badge.

Above the columns are the summary tiles **Active**, **Review**, **Done**, and **Deliverables** — a quick overview of how much work is running, what's waiting for approval, what was approved today, and how many files the team has delivered in total.

Filters and switching views

- **Filter by agent** — click an agent's name tag above the board to see only their tasks; **All** clears the filter.
- **List view** / **Kanban view** — the switch at the top left above the board. The list comes in handy as tasks pile up; on a phone the list is always used.
- In the list view you also get the tabs **Active**, **Blocked**, **Review**, **Done**, **Cancelled**, and **All**.

You can drag cards between columns with the mouse. Two exceptions: tasks in the **Review** column can't be dragged (they're waiting for a decision in the detail), and a task can't be moved to the **Done** column manually — the app reports: "This task cannot be manually moved to Done. Send it to review and approve it in the task detail; approval will finish it automatically."

Three ways a task is created

1. Via a brief in the CEO chat (recommended)

Write the brief to the CEO in the **Chat** section — they create the task, think it through, and delegate it to the right specialist themselves. Details in the Chat with Agents chapter.

2. Manually on the board

1. In the **Tasks** section, click **New Task**.
2. Fill in the **Title** and **Description** (this is where the whole brief goes — see the best practice below).
3. Choose the **Priority** (**Low** / **Medium** / **High** / **Critical**) and pick an agent in the **Assign to** field. Optionally add **Tags (comma-separated)**.
4. Attach the source materials right in the **Attachments** section — drop files there, or click to select. The same limits as in chat apply (images 5 MB, documents 20 MB, audio/video 100 MB).
5. Click **Create Task** — the task appears in the **New** column.

3. With the /task command in Telegram

Send `/task` plus the brief text to the CEO topic in Telegram — the task is created automatically and the CEO delegates it. See Communication Channels.

> ✅ **Best practice:** A good brief has four parts: the **goal** (what should be produced and why), **source materials** (links to documents in Data Center, examples, data), the **deadline**, and the **output format** (e.g. "a CSV table", "3 post variants, each max 150 words"). The more precise the brief, the fewer feedback loops.
>
> ⚠️ **Watch out:** Two common pitfalls. (1) Ten tasks at once with no priorities — the team doesn't know where to start and you drown in reviews; assign in batches and mark what's **High** priority. (2) A task without source materials — the agent fills in the missing context by guessing and you end up returning the result anyway; upload the materials to Data Center in advance and reference them in the brief.

Task detail

Clicking a card opens the detail. The header shows the number, title, status, priority, and assigned agent. For a task assigned to an agent, the detail has two tabs: **Results** and **Conversation**.

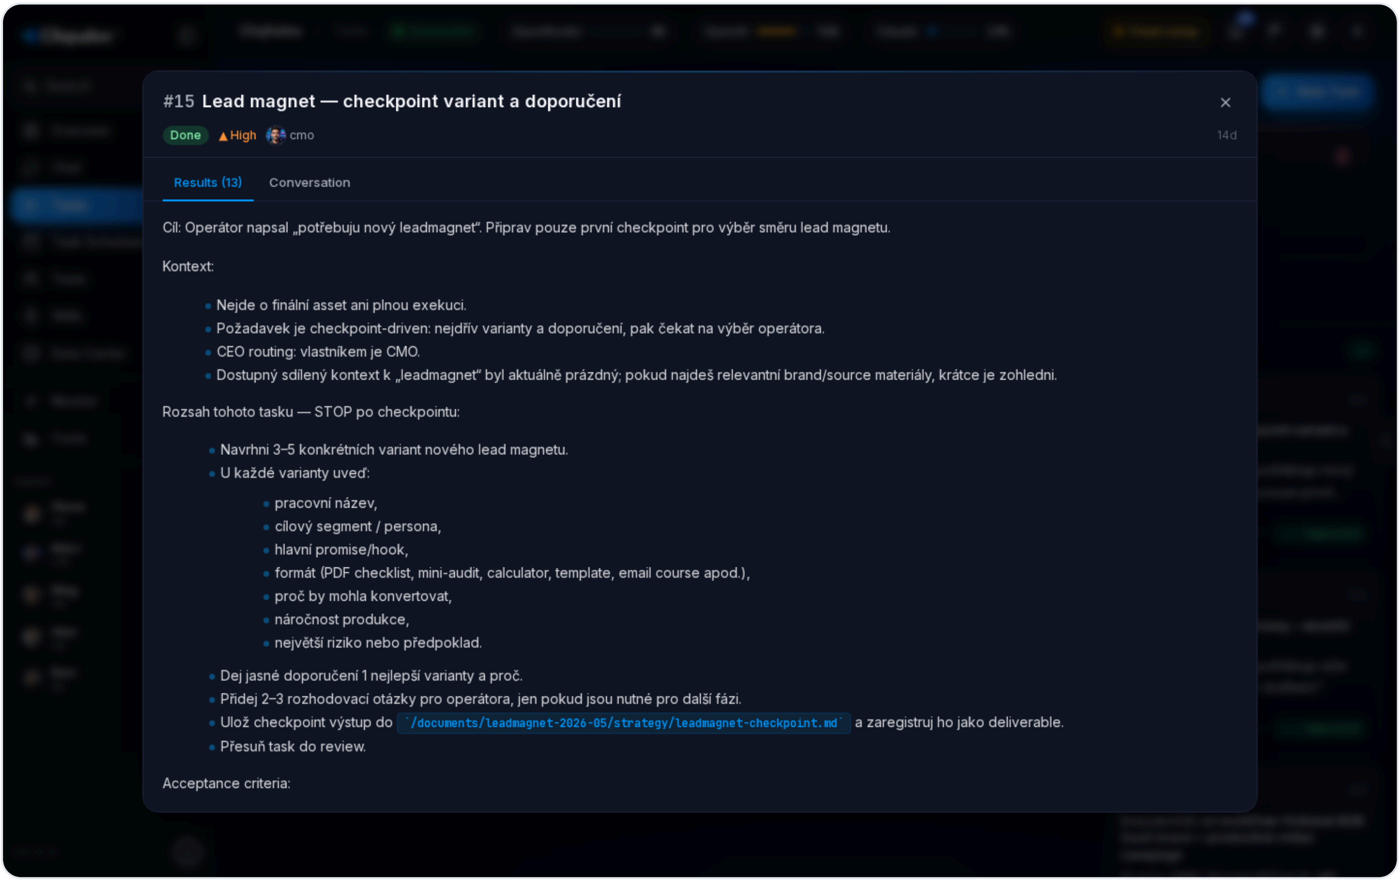

The Conversation tab — where to respond

Every task runs in its own conversation. On the **Conversation** tab you see the brief, your comments, and the agent's live replies while the work is in progress — and at the bottom there's a field to reply directly. This is where to write clarifications, additional materials, and feedback; the agent picks them up in the same thread.

> 💡 **Tip:** When a task hasn't moved for a while, open the **Conversation** first — the agent often describes there what it's waiting for (your approval, a missing document).

The Results tab — brief and outputs

On the **Results** tab you'll find the brief text, the **Deliverables** section (what the agent created), any dependencies on other tasks, the history of internal reviews, and tags.

Working with outputs:

- Clicking an output opens a **Preview** — images, PDFs, HTML, and text documents display right in the app; you can download the file (**Download**), and open links

with the **Open Link** button.
- For a task in review, the outputs are presented for assessment directly: a single output as the **Primary output**, multiple outputs as **Related outputs** with an **Expand all** option. The **Download all files** button downloads everything at once as a ZIP.
- Outputs are also stored in **Data Center → Task Outputs**, where you'll find them gathered across all tasks (see Data Center).

How to approve or return a result

When an agent finishes the work, it moves the task to the **Review** column (status **Review**, or **Quality Review**). If a task with a finished output isn't in review yet, you can send it there yourself with the **Send to review** button in the detail (it doesn't work without an output).

In the detail of a task in review you have two buttons:

- **Approve** — a window opens where you can (but don't have to) write feedback: what was good, tips for next time. With text entered, the button changes to **Approve with feedback**. Approving completes the task automatically and moves it to **Done**.
- **Rework** — you write what needs to change (required) and confirm with the **Return for rework** button. The agent is notified and reworks the task.

If the agent's run ended with warnings, the app shows them during approval in the **Agent run warnings** block. Either return the task for rework, or approve it with the **Approve despite warnings** button — with a short reason (at least 10 characters) explaining why it's okay.

A task that no longer makes sense can be ended with the **Cancel Task** button — it stays available in the list under the **Cancelled** filter.

> ✅ **Best practice:** Fast feedback = faster iterations. Don't let tasks hang in **Review** — a short "approved, shorter intro next time" takes a minute, and the team learns from every approval with feedback. The first version is rarely final; plan on one **Rework** loop and write your notes specifically ("cut paragraph 2 in half, add pricing").

Blocked tasks

When a task hits an obstacle, it gets the **Blocked** status. An orange banner "Warning: N blocked tasks" appears above the board — the **Show** button switches to a list of just the blocked tasks.

In a blocked task's detail you'll find a **Block reason** box with an explanation. Once the cause is removed (e.g. a missing document added in the **Conversation**), restart the task with the **Retry** button.

What the Failed status means

The **Failed** status means the task couldn't be completed — you'll find it in the **Done** column with a red dot. Open the detail: the **Conversation** tab shows where the work stopped. The **Retry** button gives the agent another attempt; it often helps to first add what was missing in the conversation. If you're stuck, ask the CEO in chat: "What happened with task #12?"

Where to next

- Don't keep assigning recurring tasks by hand — use the Task Scheduler.
- Assigning tasks by voice and from your phone: Communication Channels.

06 Task Scheduler

The **Task Scheduler** (in the left menu) takes care of work that should happen regularly or at a set time — without you having to assign it manually every time. A typical example is the morning briefing: set it up once and the CEO then sends you an overview of what's going on every morning, all by themselves. In the same way you can schedule a weekly report, regular content, or a sales pipeline check.

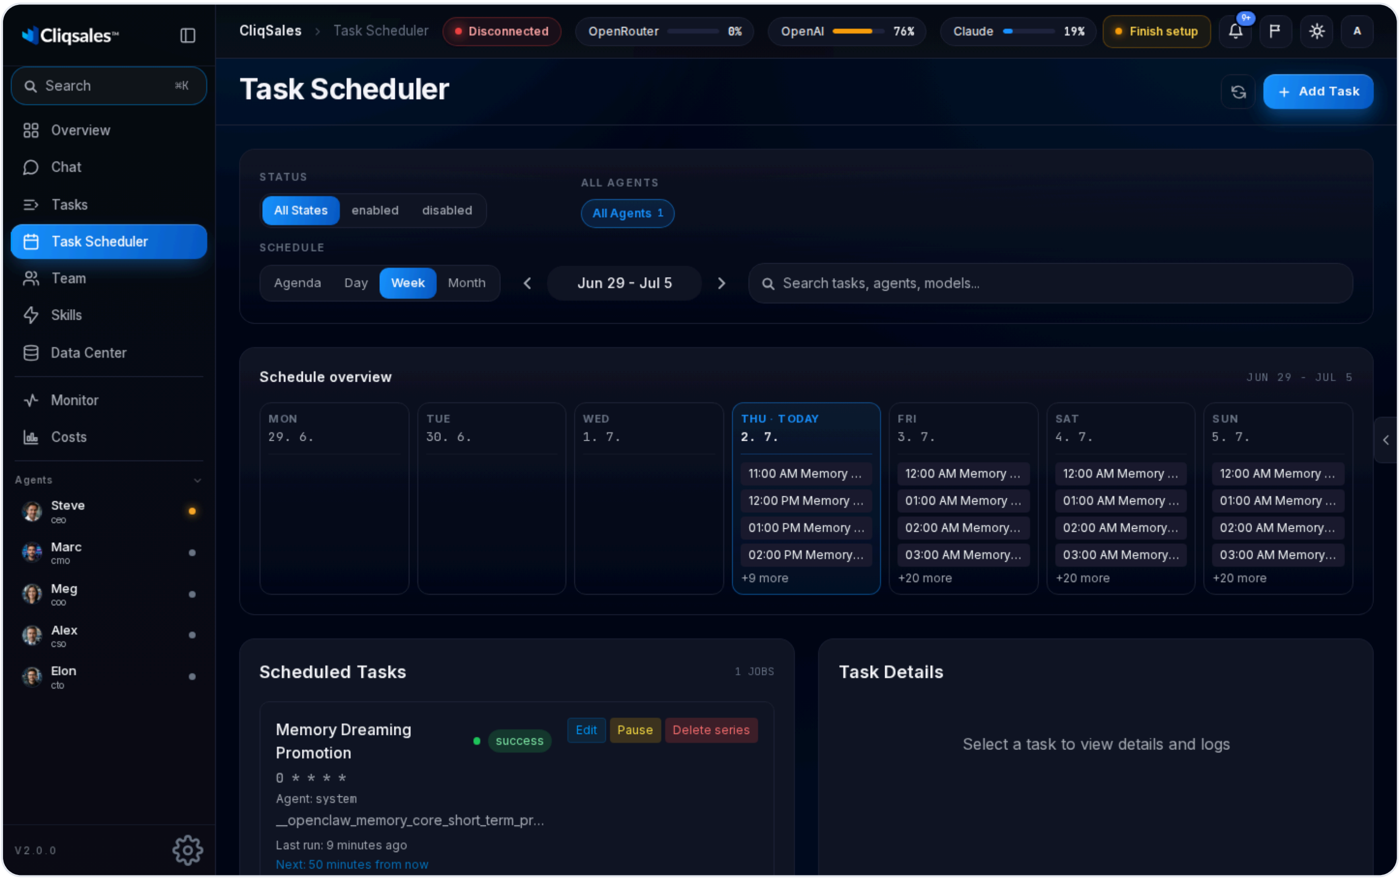

What you see in the Scheduler

- **Schedule overview** — a calendar of upcoming runs. Switch between the **Agenda** (a chronological list for the next two weeks), **Day**, **Week**, and **Month** views; browse with the arrows and return to the current date with the **Today** button.
- Filters above the calendar: by state (**All States** / **enabled** / **disabled**), by agent, and a search field ("Search tasks, agents, models…").

- **Scheduled Tasks** — the list of all jobs. For each you see the name, a colored status dot (green = enabled), the schedule, the agent, the job text, and times in natural language: **Last run** ("2 hours ago") and **Next** ("2 hours from now").
- **Task detail** — clicking a job shows its **Configuration** on the right (**Schedule**, **Agent**, **Command**, optionally **Model**, **Status** 🟢 Enabled / 🔴 Disabled, **Next run**) and **Recent Logs** — the job's run history, where you can tell whether the last run went fine.

How to create a new job

1. Click **Add Task**.
2. Fill in the **Task Name** (e.g. "weekly-report").
3. Pick the **Agent** — this determines who will handle the job (report → COO, content → CMO, overviews → CEO).
4. Click the **Schedule** field — the **Choose schedule** window opens, where you pick the recurrence without any technical knowledge: **Every day**, **Weekdays**, **Once a week**, **Once a month**, or **Selected days**, plus a **Time**. Confirm with the **Use schedule** button.
5. In the **Command** field, describe what the agent should do at the scheduled time — write it like a normal task brief: "Prepare a weekly overview of completed tasks and work in progress, and compare it with last week."
6. You can leave the **Model** field empty — the agent's default model is used. A selected model applies only to this job.
7. Save with the **Add Task** button (when editing, **Save Changes**).

In the model picker (**Choose model**), models are grouped by provider. For subscriptions you'll also see the limit status — e.g. "resets in 2h" tells you when the subscription capacity renews. Providers marked "not configured in Integrations" aren't connected; you connect them in Integrations.

> 💡 **Tip:** You don't have to create the job by hand — just ask the CEO in chat: "Every Monday at 8:00, prepare my weekly report." The CEO creates the job in the Scheduler for you.

Enabling, disabling, and editing

- **Pause** / **Resume** — temporarily turns a job off or back on (e.g. over a holiday). The job's settings are kept.

- **Edit** (in the detail, **Edit series**) — change the time, frequency, job text, agent, or model.
- **Delete series** — removes the job including all future runs; the app asks for confirmation before deleting.

The morning briefing

A typical first job to set up in the Scheduler: every morning the CEO goes through the tasks and agents and sends you a concise overview — in chat, or via Telegram (see Communication Channels).

Setting it up takes a minute: create a job (see above) with the **Every day** schedule and the time you want the overview, and use a command like: "Go through the tasks and agents and send me a concise morning overview: what was completed, what needs attention, recommended next steps, and open questions."

- **Changing the time:** find the briefing job in the **Scheduled Tasks** list, click **Edit**, and set a different time in the **Schedule** field.
- **Turning it off:** click **Pause** on the briefing job.
- **An overview on demand:** write "Give me a briefing" to an agent in chat at any time — you'll get a current overview right away, even without a scheduled job.

What's worth scheduling

- **A weekly report** — Monday morning, a summary of last week and the plan for the current one.
- **Regular content** — e.g. every Thursday, post drafts for the following week; you then just approve them in Tasks.
- **A pipeline check** — a regular report on which deals or tasks haven't moved in a while and need a nudge.

> ✅ **Best practice:** Less, but useful. Every job run consumes your subscription capacity just like a regular task — see Costs and Limits. Two or three jobs whose output you actually read are worth more than ten automated messages you delete.
>
> ⚠️ **Watch out:** Overly frequent jobs are the most common needless drain on capacity. An "every hour" check sounds tempting, but a daily summary usually tells you the same thing — for a fraction of the capacity, which would otherwise be missing for real work. When in doubt, start with a weekly frequency and increase it once the output has settled into your routine.

07 Data Center

The **Data Center** is your AI team's storage. This is where you upload the source materials the agents should work from, and where you'll also find the results of their work. It has three tabs: **Documents** (your materials), **Task Outputs** (what the team created), and **Wiki** (the company memory — it has its own chapter, Team Wiki).

Documents — source materials for your team

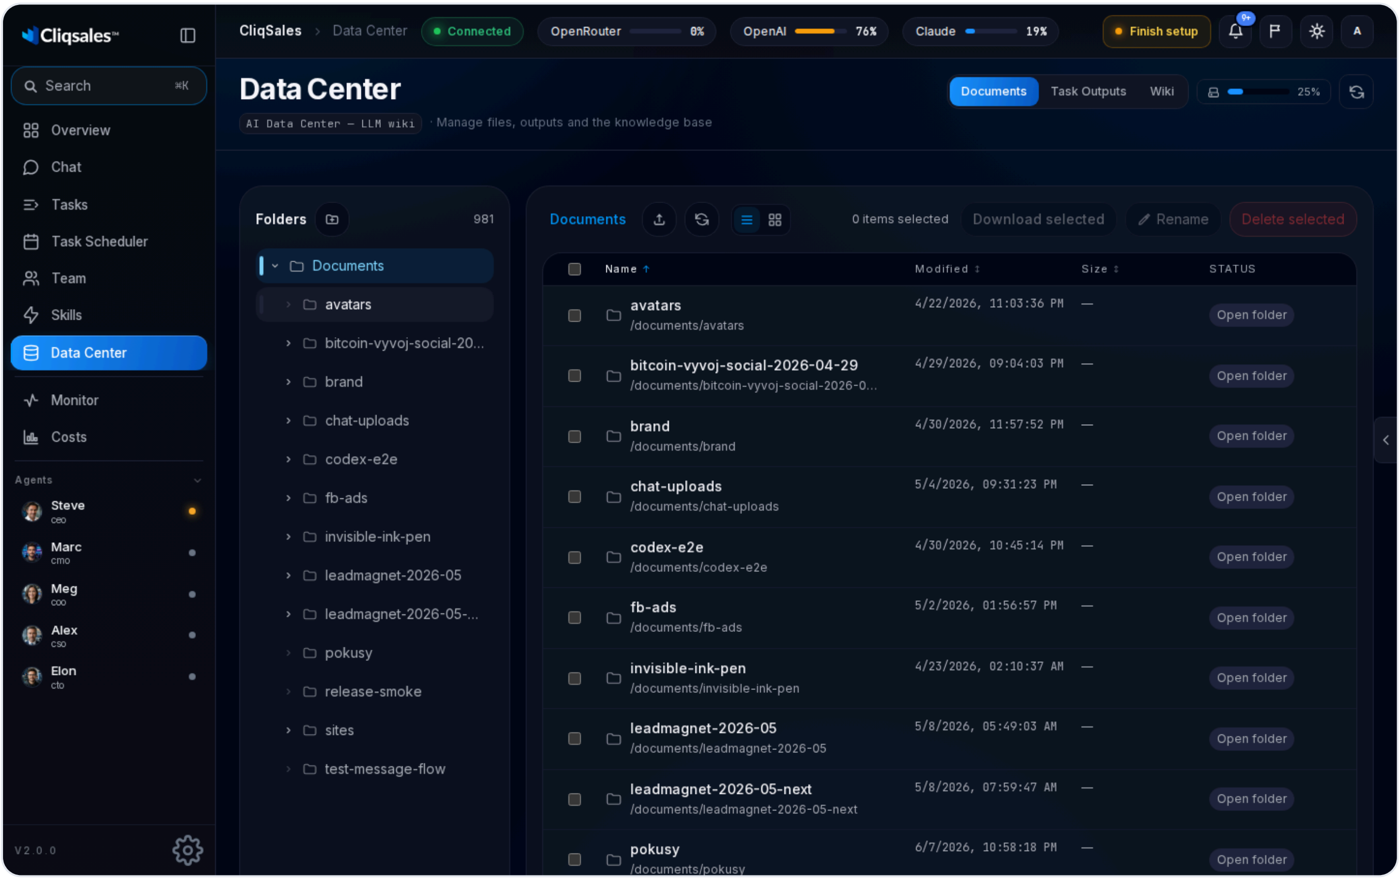

Documents is for everything the agents should draw on: the company profile, product and service descriptions, brand guidelines, templates, price lists, or data files. Agents can use uploaded documents as context while working, and the documents can also become source material for the team wiki.

On the left is the **Folders** tree, in the middle the contents of the current folder, and on the right a preview of the selected file.

> ✅ **Best practice:** Upload your core materials right at the start — the company profile, brand guidelines, templates, and product data. From the very first task, the team then works with your reality rather than generic guesses.

How to create a folder

1. Open **Data Center → Documents**.
2. Click **New folder**.
3. Enter a name — e.g. `brands/acme` or `product-01`. A slash creates a subfolder right away.
4. Confirm with the **Create folder** button.

> ✅ **Best practice:** Keep the structure simple and predictable — e.g. brand materials, product materials, and working references kept separate. Give files descriptive names ("price-list-2026.pdf", not "final_v3(2).pdf") — agents and colleagues alike will find them more easily.

How to upload documents

1. Open the target folder.
2. Click **Upload files** and select one or more files.
3. Or simply **drag the files** into the window — you can even drop them onto a specific folder in the tree. The app highlights the target and tells you where the file will go ("Drop to upload into …").
4. Follow the progress in the **Uploading files** window — each file shows its status (Waiting / Uploading / Uploaded / Not uploaded).

> 💡 **Tip:** If you upload a file that already exists in exactly the same version in the same place, Control Center won't store it twice — it just lets you know. So there's no need to worry about duplicates.

Limits and supported types

File type	Extensions	Size limit
Documents	PDF, TXT, MD, DOCX, CSV, JSON, YAML, HTML	**20 MB**
Archives	ZIP	**100 MB**
Images	PNG, JPG, GIF, WebP, SVG	**100 MB**
Video	MP4, WebM, MOV, M4V	**100 MB**
Audio	MP3, WAV, M4A, AAC, OGG, FLAC	**100 MB**

> ⚠️ **Watch out:** A file over the limit or of an unsupported type is rejected, and the app tells you why. Split large spreadsheets into smaller ones, compress long videos or upload just the relevant part. Convert exotic formats to one of the supported ones (ideally PDF or text).

File previews

Clicking a file opens a preview on the right:

- **PDFs, images, audio, and video** play or display right in the app.
- **HTML** is shown as a safe preview.
- **Text formats** (TXT, MD, CSV…) are shown as text.

The **Status** column of a document tells you whether text could be prepared from it for the agents: **Ready** means the agents can read the content; **Stored only** means the file is safely stored, but no text is extracted from it.

> 💡 **Tip:** Notes in Markdown format (`.md`) can be edited right in the preview with the **Edit** button — no need to download the file and upload it again.

How to work with a ZIP archive

1. Click the ZIP file — the preview shows the archive contents: the number of entries, files and folders, the size, and the extraction target. Nothing is extracted yet.
2. Click **Extract** — the files are extracted into a folder next to the archive.
3. System junk from macOS (`__MACOSX`, `.DS_Store`, and the like) is ignored automatically — it won't clutter your documents.

> ⚠️ **Watch out:** An archive with unsafe paths or contents that are too large can't be extracted — the app says so next to the **Extract** button. In that case, extract the archive on your own computer and upload only the files you need.

Gallery view

The **List** / **Gallery** switch at the top right changes how the folder is displayed. **Gallery** shows images as thumbnails — handy for logos, banners, and other visual materials where you need to see the content at a glance.

How much space you're using

The Data Center header has a **Disk** indicator — the percentage of server space in use. When it's getting full, delete files you no longer need — you can do it in bulk via **Delete selected** (the wiki evidence is kept).

How to check a documents folder

Control Center keeps a registry of documents (what was uploaded, from where, and in which version). A folder check compares the registry with the folder's actual contents:

1. Open the folder you want to check.
2. Click **Check folder**.
3. The **Document evidence check** result sorts the findings into groups:

- **Files without evidence** — new files the registry doesn't know yet.
- **Evidence without a file** — records whose file is no longer in the folder.
- **Changed content** — files whose content differs from the registry.
- **Skipped** — files skipped e.g. due to size or type.

1. Resolve the findings with the **Register found files** button (adds new files to the registry) and **Mark missing as removed**.

The check itself deletes nothing and doesn't display document contents. Registering files does not automatically add them to the wiki.

Renaming, downloading, and deleting

Select items with the checkboxes in the left column; actions then appear above the list:

- **Rename** — changes only the file or folder name. Wiki evidence and audit history stay linked to the same source.
- **Download selected** — downloads the selected items as a single ZIP bundle.
- **Delete selected** — deletes the selected items after confirmation. Wiki evidence and audit history are kept — deleting a file doesn't erase what the team has already learned from it.

> ⚠️ **Watch out:** Rename and delete files via the actions in Data Center, not outside the app. When a file disappears some other way, the folder check flags it as missing and Control Center doesn't know what happened to it.

Task Outputs — what the team created

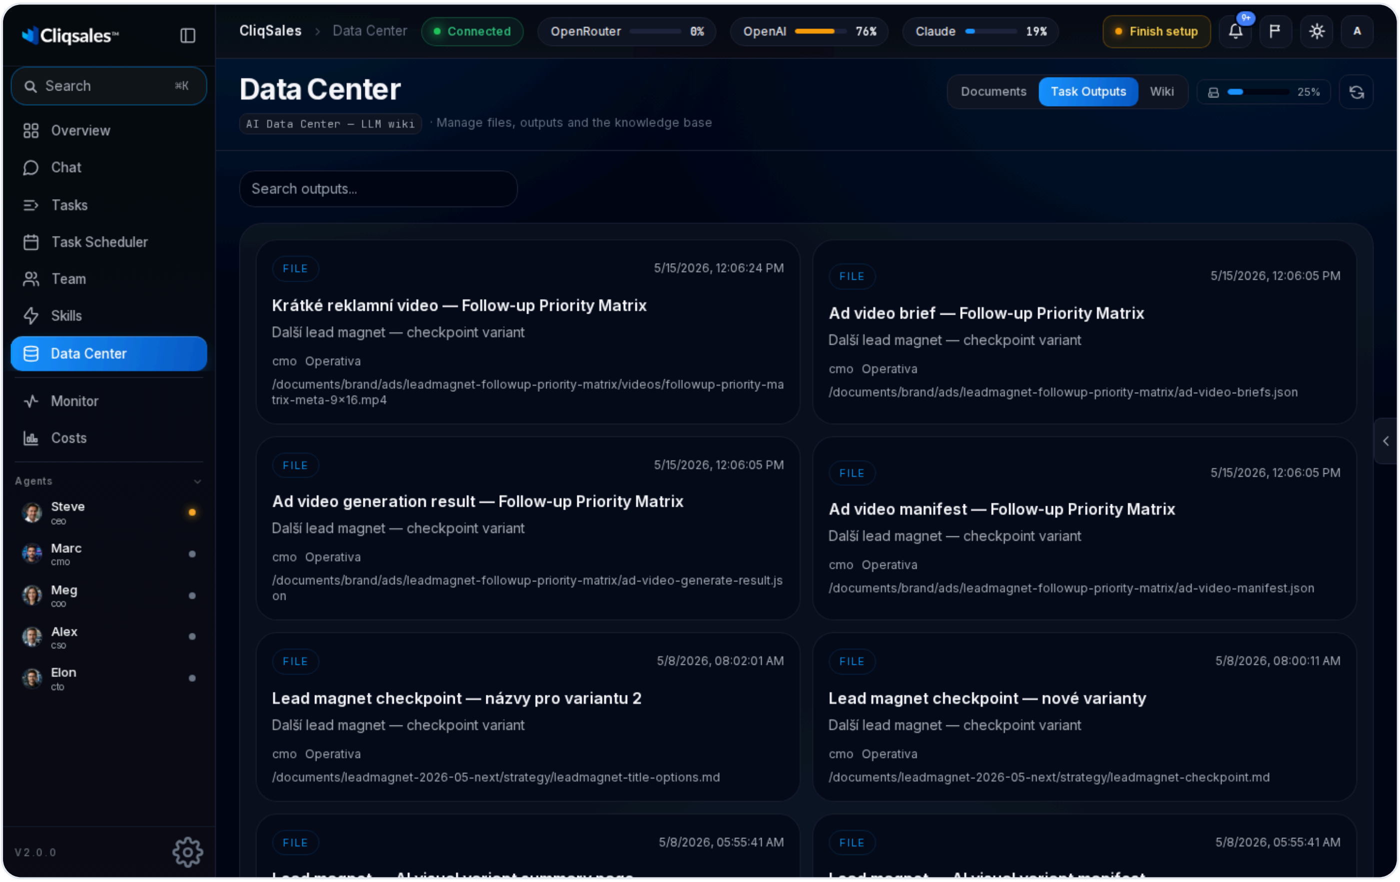

The **Task Outputs** tab automatically collects the outputs (deliverables) the team saves while working on tasks and projects — texts, reports, links, and files. The same outputs are also in the relevant task's detail (see Tasks); here they're all in one place.

How to find and download an output

1. Open **Data Center → Task Outputs**.

2. In the **Search outputs…** field, type a name, task, project, or author.
3. Click an output's card — the detail opens with the type (**Text**, **Link**, **File**, **Report**), the task, the author, and the date.
4. The detail shows a preview right away: text, image, PDF, HTML, and even ZIP contents. For links, use **Open Link**.
5. The **Download all files** button downloads the output's files as a single ZIP bundle.

Until the team saves its first output, the tab is empty — outputs appear with the first completed tasks.

Wiki

The third tab, **Wiki**, is the team wiki — the company memory Control Center builds from selected sources. In the Data Center you read and manage it; how it works and how to set it up is covered in its own chapter, Team Wiki.

08 Team Wiki

The team wiki is your AI team's company memory. From selected sources — meetings, documents, emails, and task outputs — it automatically builds clear, searchable pages the team can return to at any time. You'll find it in **Data Center → Wiki**; it has four parts: **Browse**, **Graph**, **Jobs**, and **Settings**.

What the wiki is for — and what it doesn't do

The wiki stores knowledge with long-term value: meeting transcripts, important documents, task outputs, selected emails, and regular overviews. It helps the team find connections across meetings, documents, and projects — and for every piece of information, its origin is traceable.

The wiki is not a chat and doesn't run the team for you. On its own it:

- doesn't send emails or publish articles,
- doesn't create new tasks,
- doesn't delete original documents,
- doesn't read your entire mailbox,
- doesn't automatically give agents the whole company memory.

For selected jobs, agents can receive just a short selection of relevant information from the wiki — and only when you enable it in **Settings** (see below).

> ✅ **Best practice:** Feed the wiki business and strategy meetings, client decisions, workshop outputs, important project reports, and documents the team will come back to.
>
> ⚠️ **Watch out:** Newsletters, automated notifications, personal mail, duplicate transcripts, and one-off technical records don't belong in the wiki. The wiki works best as a curated company memory, not a dumping ground for everything.

How wiki cards are created

When the wiki receives new material, it first stores it as a traceable source and then builds clearer pages from it. A page can include a summary, important decisions, the people and companies mentioned, projects, open questions, or notes about what still needs verifying.

The wiki content is organized into five groups (you'll see them in the tree in **Browse**):

- **Cards** — the specific people, companies, projects, and products the wiki recognizes in the sources.
- **Topics** — more general concepts and recurring themes.
- **Summaries** — broader summaries connecting multiple pieces of information.
- **Materials** — the original sources (a meeting, document, email…) the pages were built from.
- **Alerts** — generated overviews and reports, for example the daily overview.

The basic card types (**Company**, **Person**, **Project**, **Product**) are fixed by the system so the wiki stays tidy. New specific items (a new company, a new project) are created by the wiki itself as soon as they appear in a source.

> 💡 **Tip:** With the **New type** button in **Browse** you can add your own card type (e.g. "Property"), including a description of what the wiki should track for it. Deleting a type, however, also deletes its linked cards — think it through first.

Browse — reading the wiki

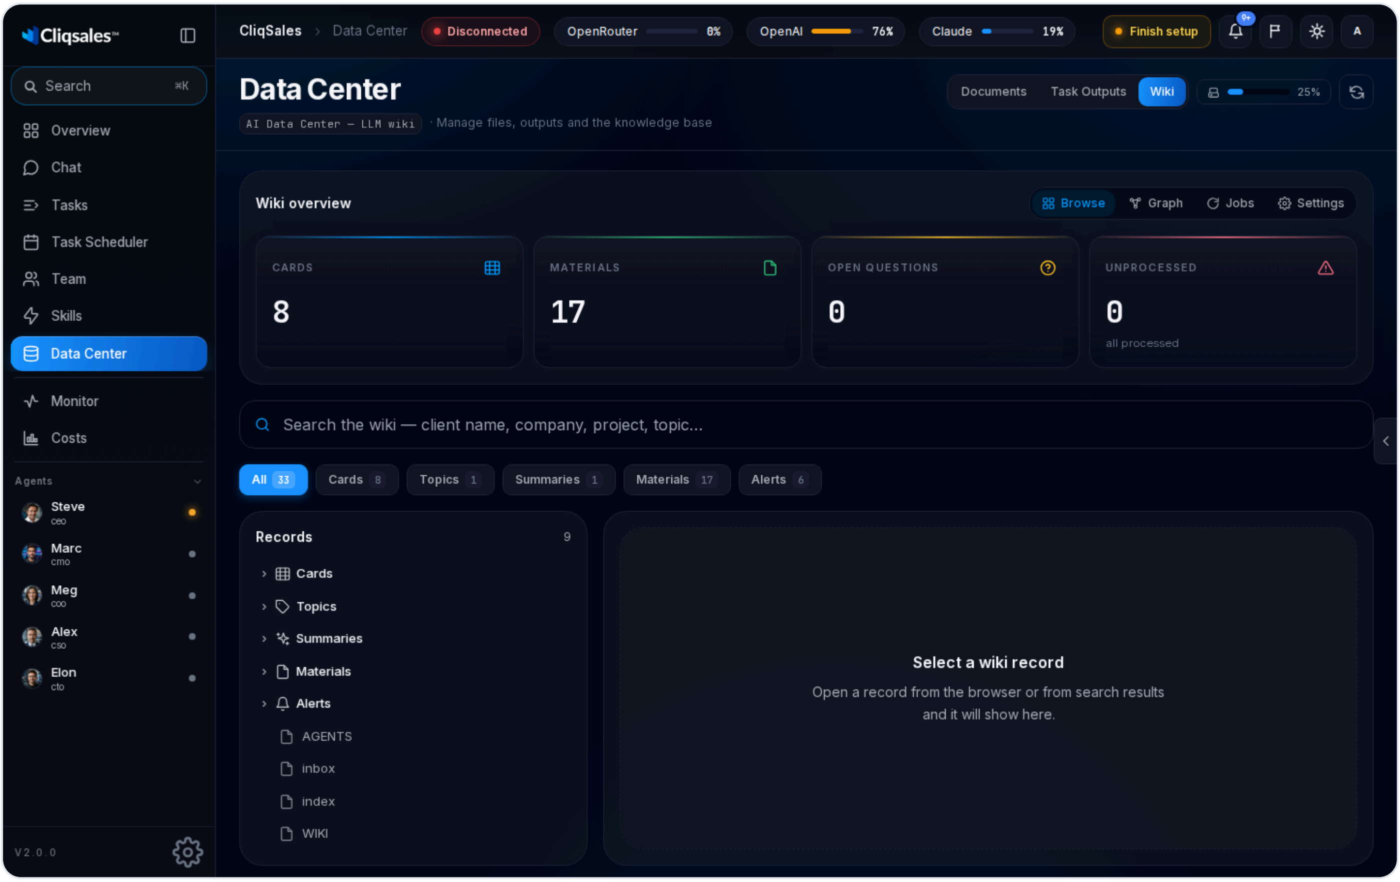

Browse is for reading. At the top is the **Wiki overview** with counts: **Cards**, **Materials**, **Open questions**, and **Unprocessed** (how many materials are still waiting to be processed). On the left is the page tree — click to open a folder or page, and the content appears on the right.

How to find a page

1. In the **Search the wiki — client name, company, project, topic...** field, type what you're looking for.
2. Results appear in the left panel; click a page to open it.
3. You can narrow the card list with the type filter (**Companies**, **People**, **Projects**, **Products**) and the state filter (**Clients**, **Leads**, **Suppliers**, **Partners**, **Archive**).

What's on a card

A card summarizes what the team knows about a given person, company, or project: the **What we know** section, **Open questions**, **Source materials** (where the information came from), and **Related**.

> 💡 **Tip:** You can answer open questions right on the card — the **Fill in** button opens a field for your answer. That's how you enrich the wiki with things only you know.

Graph — a map of connections

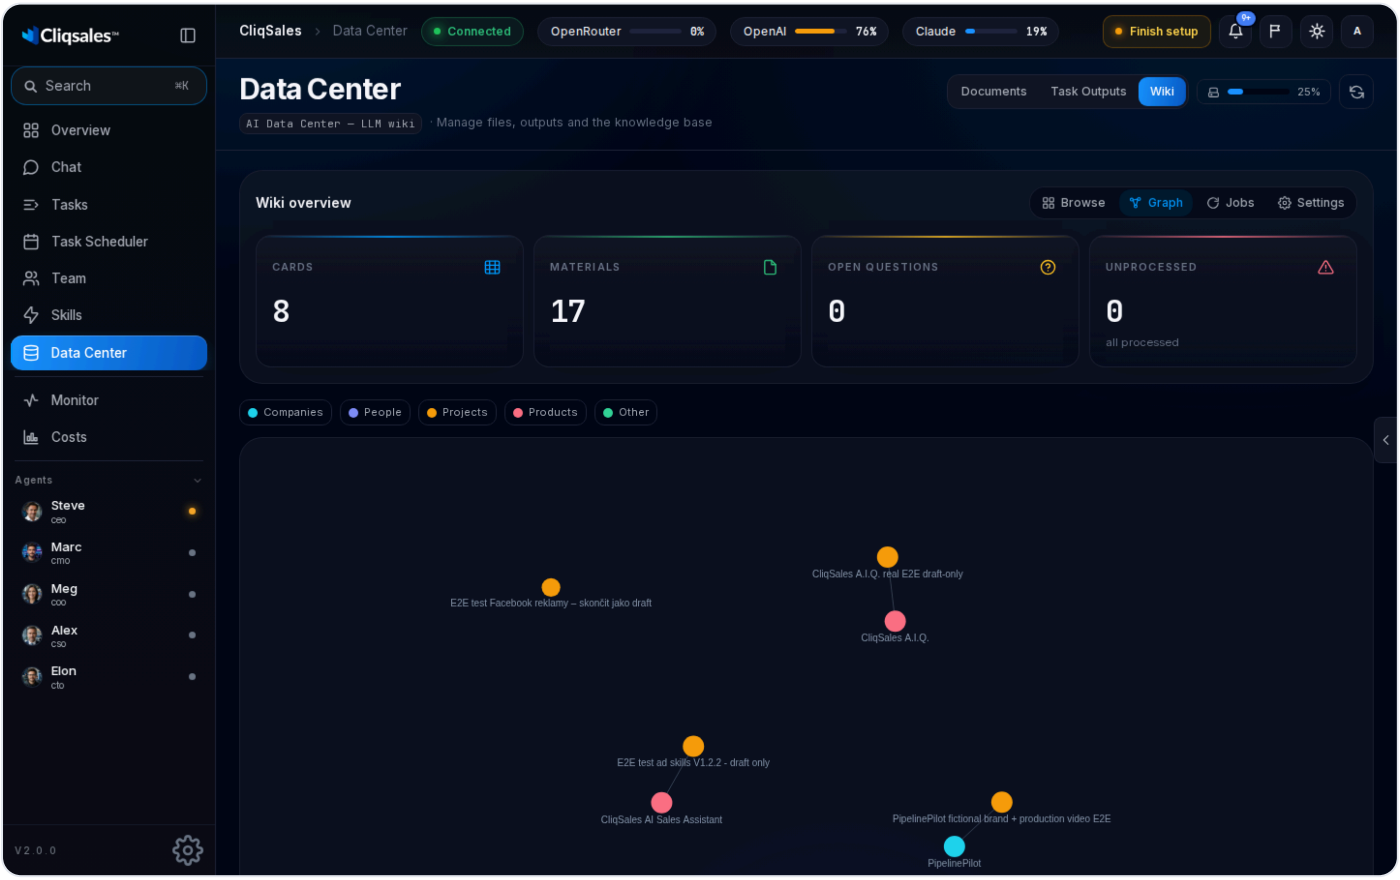

The **Graph** shows the relationships between wiki cards as a visual map:

- Each circle is one card (a person, company, project...).
- The **color** matches the card type, the **size** matches the number of links.
- Lines show that cards are related.
- The type legend at the top (**Companies**, **People**, **Projects**, **Products**, **Other**) doubles as a filter — click a type to hide or show it.
- Clicking a circle opens the card in **Browse**.

The graph fills in on its own as the wiki processes sources. When it's empty, you'll see the message "The graph lights up once more cards are connected." — the wiki doesn't have enough connected cards yet; just add more sources.

Jobs — what the wiki received and how it went

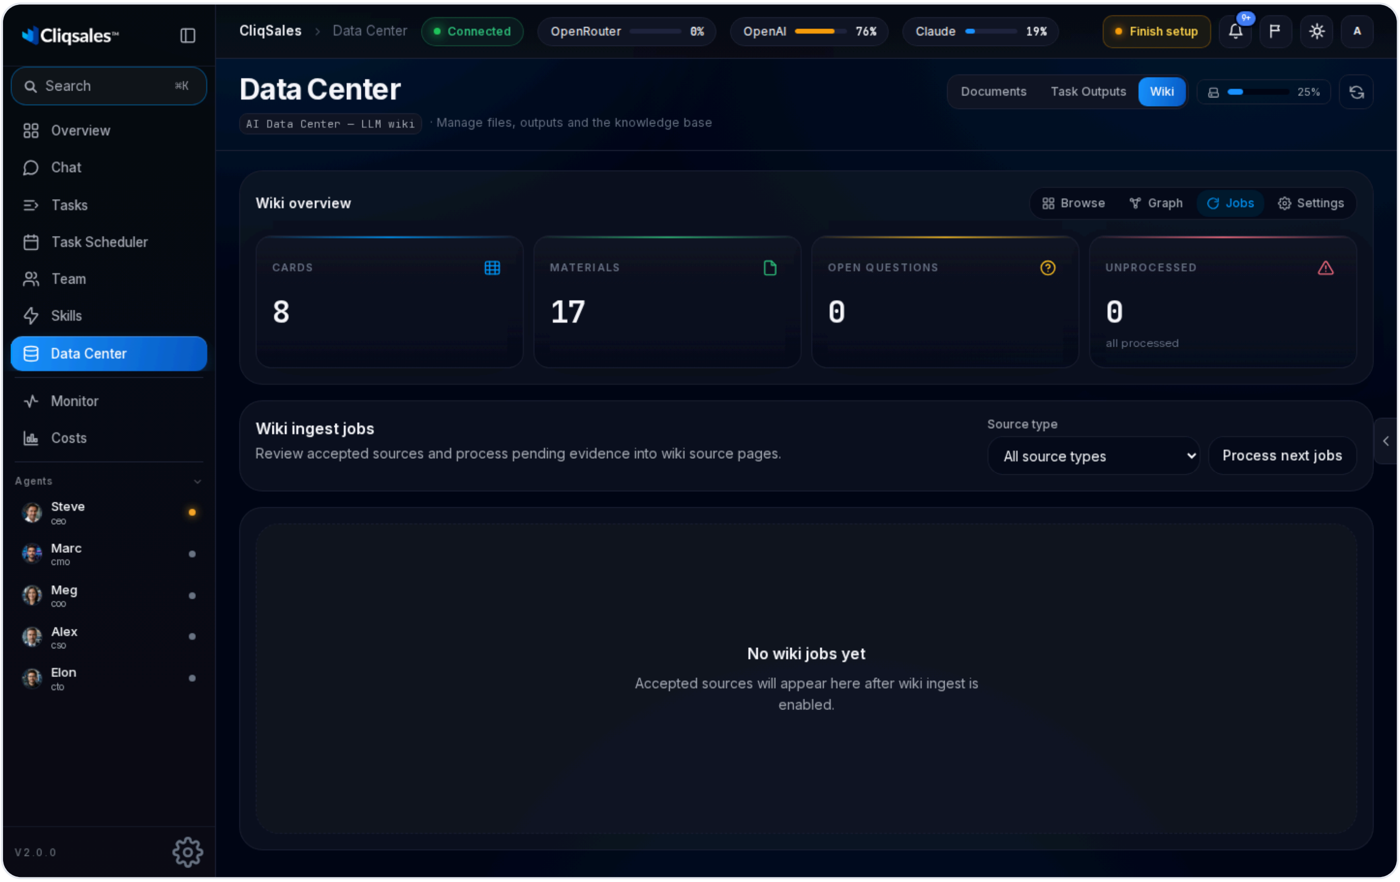

Jobs is the wiki's diary: each row is one received piece of material — an uploaded document, a Fireflies meeting, a job output, an email, or a manually added source. For each row you see the status, the source type, the time of the last change, the number of attempts, any error, and a link to the resulting wiki page.

At the top you can filter by source type (**Source type**: **Uploaded documents**, **Task outputs**, **Completed tasks**, **Email**, **Fireflies**, **Manual sources**). The **Process next jobs** button manually kicks off processing of the waiting materials.

Processing statuses

Status	What it means
Pending	The material is waiting for its turn.
Processing	Control Center is working on the material right now.
Classified	The material has been understood and prepared for writing.
Written	The result is stored in the wiki.
Ignored	The material was received, but per the rules it shouldn't be used.
Failed	Processing didn't succeed.
Dead letter	The material failed repeatedly or can't be processed safely.
Skipped (needs attention)	Automatic processing failed repeatedly (e.g. the material is too large) — it needs your attention.

The timeline

The **Show timeline** button on a record opens a step-by-step history: what happened to the source and where it may have stopped. It's useful when you're waiting for a new page and can't see it, when processing ended in an error, or when you want to find out where a page came from.

Settings — what the wiki may receive

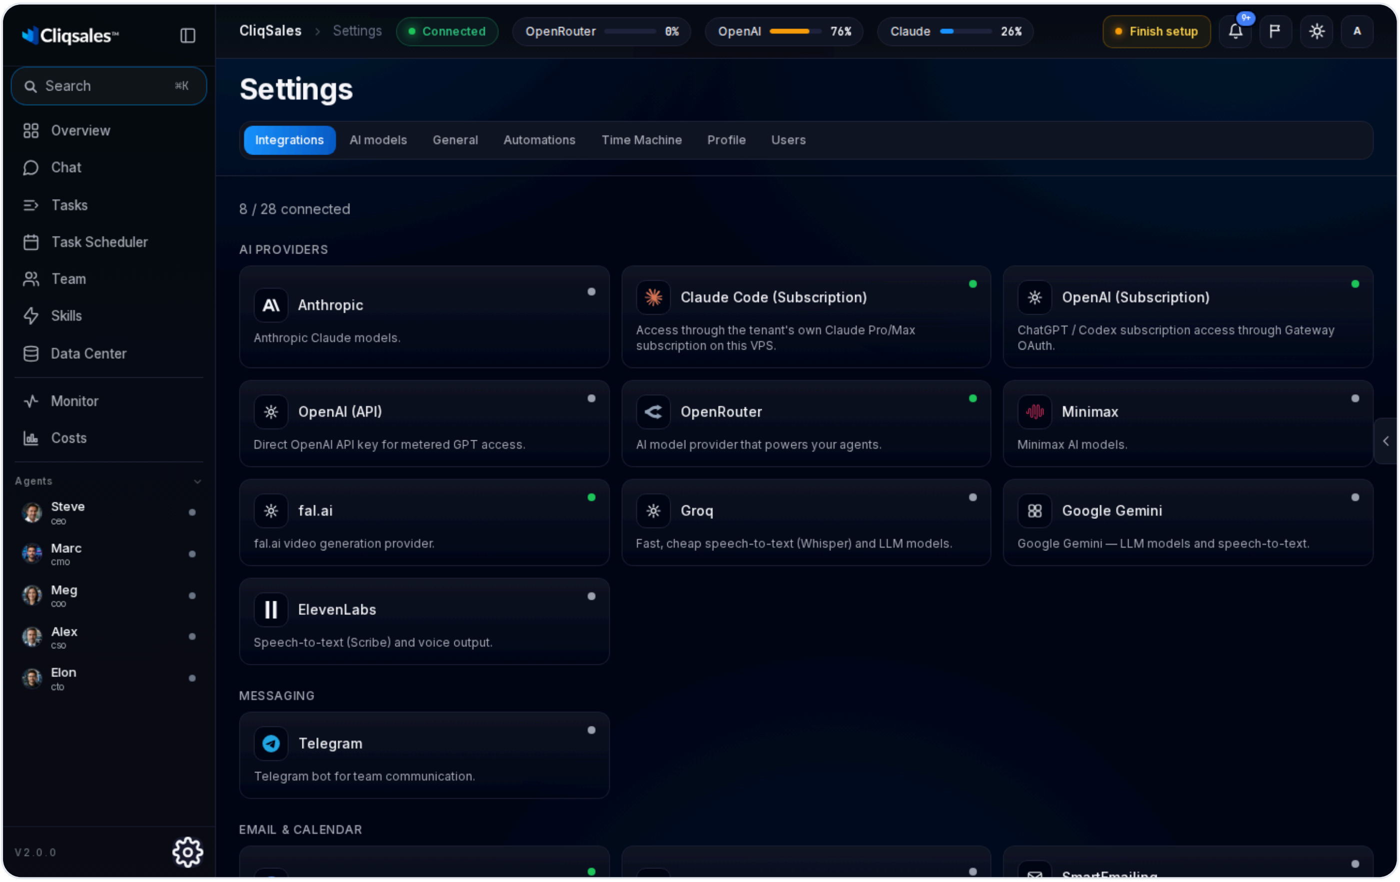

Settings determines what the wiki is allowed to receive. A regular user can read the settings; saving changes usually requires higher permissions. You confirm all changes on this screen with the single **Save all** button in the bottom bar.

The main switch — a safe pause

The **Enable wiki ingest** toggle controls whether the wiki accepts new material. When you switch it off:

- new materials are not accepted,
- the existing wiki stays visible,
- nothing is deleted automatically,
- the processing history remains available.

So it's a safe pause, not a wiki wipe.

A short wiki context for agents

The **Let agents use wiki context** toggle controls whether selected jobs can receive a short selection of relevant information from the wiki and documents before they run. When it's on:

- Control Center picks only a few related materials,
- the agent doesn't get the whole wiki,
- the selection is traceable in the job or run detail,
- if the selection fails, the job doesn't stop because of it.

Turn it on when you want agents to build better on company knowledge. If you prefer to keep the agents' work as separate from the wiki as possible, leave it off. Regular chat isn't automatically affected by this setting.

Creating cards from emails

The **Creating cards from emails** option says what happens when an email arrives from a new person: **Create automatically** (a new contact becomes a card right away) or **Suggest first** (the card is only proposed and you confirm it in the **Card proposals** section with the **Confirm** / **Dismiss** buttons). Meeting recordings always create a card; newsletters and no-reply addresses never do.

> 💡 **Tip:** In **Settings** you'll also find **Download wiki snapshot** — it downloads the whole wiki as plain markdown files for local browsing (e.g. in Obsidian). Your company memory isn't locked inside the app.

Wiki sources

Fireflies — meeting transcripts

Fireflies is a meeting transcription tool. The wiki can turn a meeting into a page; it needs a working Fireflies integration and an inbound address (the UI calls it a **webhook** — simply an address Fireflies sends a message to when there's a new meeting).

Step by step:

1. In **Settings** → **Integrations**, verify the Fireflies API key (see Integrations).
2. Open **Data Center** → **Wiki** → **Settings**, the **Fireflies** section.
3. Turn on **Enable Fireflies sidecar ingest**.
4. Next to **Fireflies webhook**, click **Create webhook** — a secured inbound address is created and the secret signing key is shown right away — copy it with the **Copy key** button. Save it immediately; for security reasons it's never shown in full later.
5. Copy the address with the **Copy** button.

6. In Fireflies, set the **Webhook URL** to the copied address. If Fireflies offers a **Signing Secret** field, paste the secret key from Control Center into it.
7. Confirm with the **Save all** button.

The wiki only accepts a secured inbound address — if you pick an existing webhook without verification, the app warns you that a webhook secured with a Bearer token or HMAC is required (two ways services verify that a message really came from Fireflies).

> 💡 **Tip:** You don't have to cancel your existing Fireflies automation. The wiki behaves as an add-on — it only takes a copy of the received event for the company memory.

Email — strictly selective

The wiki shouldn't crawl your whole mailbox by itself — email contains private messages, notifications, and a lot of noise. Use the selective setup: a chosen folder in your mailbox or chosen Gmail labels. Until you pick a mode in the **Email** section, no email flows into the wiki; the recommended mode is **Selective** (only emails from the sources you add yourself are processed).

Recommended approach:

1. In your email, create a folder, for example **Company Wiki**, and move only the messages that truly belong in the company memory into it.
2. In Control Center, add the email account in **Settings → Integrations** (see Integrations).
3. Open **Data Center → Wiki → Settings**, the **Email** section, and choose the **Selective** mode.
4. Under **Email sources for wiki**, click **Add source**: pick the **Source type** (**IMAP account** or **Google Workspace Gmail**), the account, and the **Folder**, or the **Gmail labels**.
5. Click **Test without writing** — the test only shows which messages the wiki would find (Accepted / Ignored / Duplicate); nothing is written.
6. When the result looks right, run **Check now** — the accepted messages are queued for processing. Follow the progress in **Jobs**.

During a regular check, the wiki doesn't delete or move messages, and doesn't mark them as read.

> ⚠️ **Watch out:** The **Full** (beta) mode scans the complete mail of the selected accounts. That makes it easy to let newsletters, notifications, and personal messages into the company memory. Stick with the **Selective** mode with a folder or labels.
>
> 💡 **Tip:** The **Exclude Rules** section can permanently filter out specific senders, domains, or keywords — useful for the recurring noise that occasionally sneaks into your selected folder.

Documents from Data Center

Documents uploaded to **Data Center → Documents** can be wiki source material — including files extracted from ZIP archives and documents renamed via the app's actions. Details on uploading and the folder check are in the Data Center chapter.

Task outputs

Approved or saved task outputs can be wiki source material — important conclusions, research, and reports thus make it into the company memory. Not every passing detail from chat belongs in the wiki; finished outputs are what matter most. More in the Tasks chapter.

The daily overview

The daily overview is an optional summary of what changed in the wiki over the previous day: received sources, open questions, and processing warnings. Control Center writes it once a day as a report into the wiki and delivers the summary straight to you — it doesn't wake an agent for it and doesn't draw on your subscription capacity.

You choose where the overview goes with the delivery toggles:

- **Send to chat** — it appears with the coordinator (CEO) as the permanent **Daily Pulse** thread.
- **Send to Telegram** — it arrives in the operator group, in the coordinator's section (see Communication channels).
- **Send by email** — pick the **From account** (you need a connected email account — see Integrations) and fill in **Send to** (comma-separated addresses). The email is sent only after the report has been stored in the wiki.

In the **Daily overview** section in **Settings** you also set:

1. **Enable daily overview** — the main toggle.
2. **Time** and **Time zone** — when the overview should be prepared.
3. **Length** — **Short**, **Standard**, or **Detailed**.
4. Save with the **Save all** button.

The **Generate now** button is for a manual test — no need to wait until the next day. You'll find the generated reports in **Browse** under the **Alerts** group; in **Daily overview history**, open an older report by clicking its record (**Open report**).

Recommended start — begin small

1. Open **Data Center** → **Wiki** → **Settings**.
2. Turn on **Enable wiki ingest**.
3. Pick one safe source for the first test — for example a single Fireflies meeting or a single document.
4. Let the source go through processing.
5. In **Jobs**, verify that the status ended as **Written**.
6. In **Browse**, find the resulting page.
7. Only then add more sources.

> ✅ **Best practice:** Start small. The wiki will be more useful when you feed it quality material, not everything that exists in the system.

When something doesn't work

I can't see a new page

1. Open **Data Center** → **Wiki** → **Jobs** and find the source.
2. Check its status.
3. If the status is **Pending**, click **Process next jobs**.
4. If the status is **Failed**, open the timeline (**Show timeline**) and look at the last error.

A Fireflies meeting didn't show up

Check, in order:

- the Fireflies API key in **Settings** → **Integrations**,
- the inbound address (**Fireflies webhook**) in the wiki settings,

- the secret key configured in Fireflies,
- the **Enable Fireflies sidecar ingest** toggle,
- the **Jobs** view, in case there's an error.

An email didn't come in

Check:

- that the email account is added in **Settings → Integrations**,
- that you selected the right folder or Gmail labels,
- that there's a test message in the folder,
- that the email source is saved and enabled in the wiki settings,
- the result of the **Test without writing** run.

The daily overview wasn't created

Check:

- that **Enable daily overview** is on,
- the configured **Time** and **Time zone**,
- **Daily overview history** — the status of the last run.

In short

The wiki is the place for company knowledge with long-term value. It gives the best results when you feed it selected, quality material and let Control Center turn it into traceable pages. Don't use it as a dumping ground for everything — use it as a company memory the team can safely return to.

09 Skills — your agents' abilities

Skills are your agents' abilities — playbooks and connections that let them do specific things: write SEO articles, publish to social media, generate product photos, and more. In this chapter you'll learn where to find skills, how to enable them for agents, and how to create your own.

What types of skills exist

- **CliqSales skills** — we ship them as part of the platform. They carry the **Distributed** badge, are **Read only**, and may change or grow after a platform update (you'll spot additions by the **New** badge).
- **User skills** — your own. They carry the **Custom** badge; you can upload them, edit them in the editor, or delete them, and several agents can share one at the same time.
- **Built-in capabilities** — some abilities (for example working with a browser or generating images) are built directly into the core your agents run on. They don't appear in the **Skills** library and are configured by the platform administrator — in day-to-day use you don't need to worry about them.

The Skills library

Open **Skills** in the left menu. There are two tabs at the top:

- **CliqSales Skills** — an overview of the distributed skills. The **Preview** button shows a skill's content. They can't be edited — the skill's source text stays in English, while the labels in the app are localized.
- **User Skills** — your own skills. This is where you upload, edit, and delete them.

Below the header you'll see a summary (**skills total**, **last updated**) and skills grouped by category — they start collapsed, so expand only the ones you want to work with.

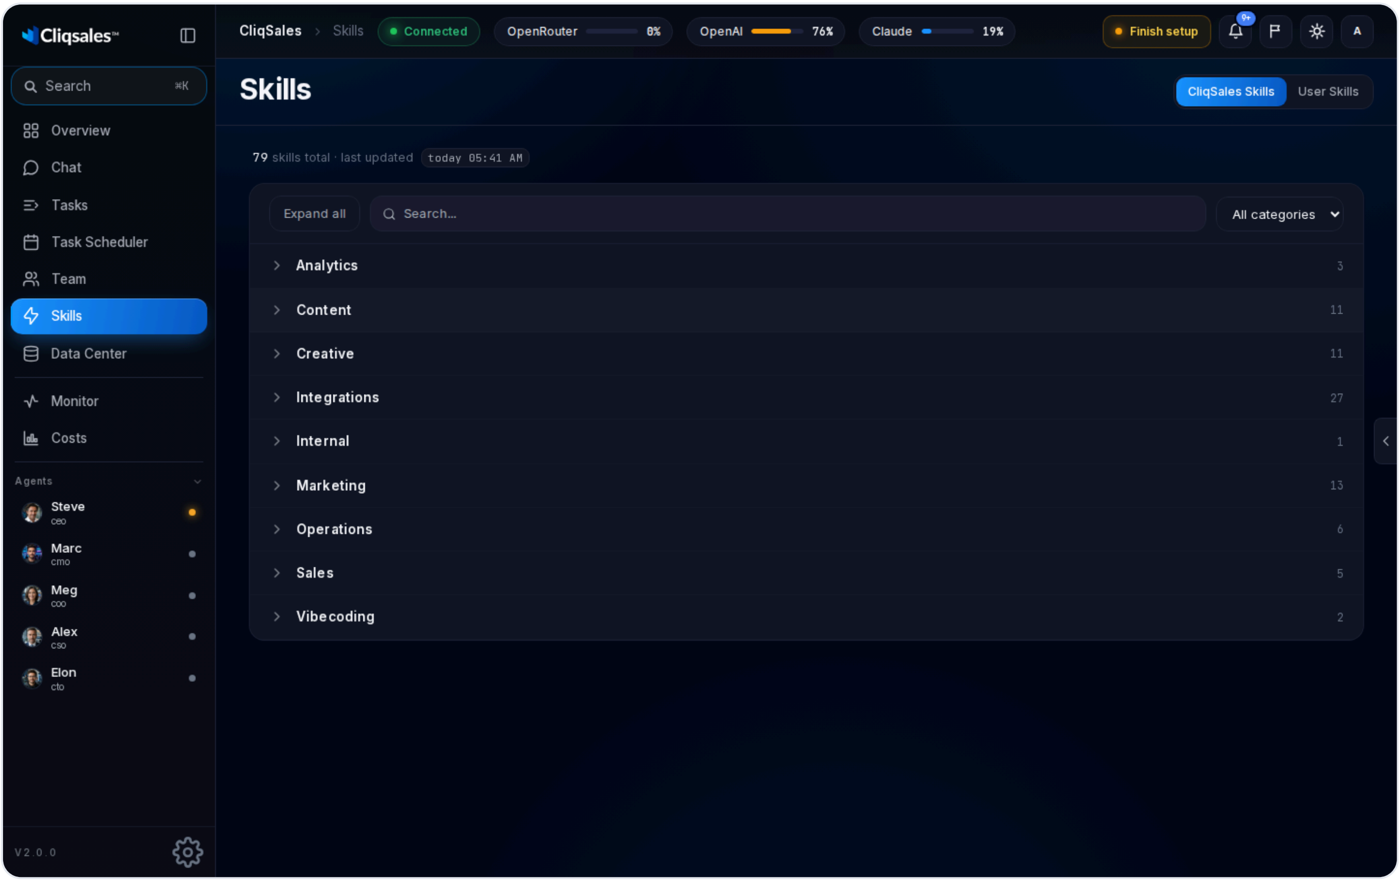

How to upload your own skill

1. Open **Skills** → the **User Skills** tab.
2. Click **Upload Skills** and use **Choose file** to pick a `.md`, `.zip`, `.tar.gz`, or `.tgz` file. The category is optional.
3. If the archive contains several skills, the app shows a list (*Found X skills — select to import*) — tick the ones you want and confirm with **Import selected**.

To change an existing custom skill, click **Edit**: you can update the **Name**, **Category**, **Description**, and the **SKILL.md content** itself, then save with **Save**.

> 💡 **Tip:** You don't have to write a skill yourself — ask the CTO agent to create it. For the result to fit on the first try, describe four things in your request: the **input** (what the agent receives), the **procedure** (the steps to follow), the **output** (what the result should look like), and an **example**. For instance: "Create a skill called *Meeting minutes*: the input is a meeting transcript, the procedure is to summarize decisions and action items with owners, the output is a structured one-page summary — I'm attaching an example of what it should look like."
>
> ⚠️ **Watch out:** Deleting a custom skill (**Delete**) removes the file from all agents that had it enabled, too. The app asks for confirmation once more before deleting.

How to enable a skill for a specific agent

Skills are enabled in each agent's detail, not in the library:

1. Open **Team** → click an agent → the **Skills** tab (the agent detail is described in the chapter Your AI team).
2. Switch between the **CliqSales Skills** and **User Skills** sections. Each shows an **Enabled** list and an **Available to enable** list; the **Filter skills…** field offers quick filtering.
3. Next to a skill, click **Assign** (enable) or **Remove** (disable).

In the **User Skills** section you can also upload a new custom skill straight from the agent detail (**Upload Skills**) or edit an existing one (**Edit**).

> ✅ **Best practice:** A few well-chosen skills beat "everything for everyone". Enable only the skills that match an agent's role — the agent then makes better decisions about which procedure to use, and the results are more predictable.

When a change takes effect

Changes are saved immediately, but they don't cut into work in progress:

- a running task is allowed to finish,
- the change takes effect on the agent's next wake-up (when it picks up its next piece of work),
- or after a session reset (a session is an agent's working block).

Disabling a skill therefore never interrupts the work an agent is doing right now.

What the Dummy and Spec Only statuses mean

Some skills show a status badge:

- **Ready** — the skill is fully functional.
- **Dummy** — the skill already has its instructions and workflow, but no real integration or built-in capability is wired up yet. It's not an empty skill; it just doesn't "pull real levers" yet.
- **Spec Only** — an even lighter variant: a draft or specification without full operational use.

Sharing a custom skill between agents

One custom skill can be used by several agents at once — it's shared. When you edit it, the change propagates to **all** agents that have it enabled.

> 💡 **Tip:** Need a special variant for just one agent? Don't edit the shared skill — create a new custom skill (for example a copy with adjusted content) and enable it only for that one agent.

Frequently asked questions

- **Can I edit or delete a distributed (CliqSales) skill?** No, they are read-only. You can, however, disable them for individual agents at any time.
- **Where do I find the core's built-in capabilities?** They don't appear in the **Skills** library or in the agent detail — the platform administrator configures them.
- **If I disable a skill, does it interrupt the agent's running work?** No — the change takes effect on the agent's next wake-up or after a session reset.

10 Communication channels

The Control Center web app is your control room — you see tasks, outputs, and costs there. But you don't have to sit at a computer to talk to your AI team: through **Telegram** you assign work and get replies from your phone, on the road, from anywhere. Agents reply just like they do in the Chat section.

Telegram

Telegram is the most capable channel. You have two options there: a **team group** with four sections (Telegram calls them "topics") and **direct messages** to the bot.

How the team group works

The group has four sections. Write only in the **CEO** section — the others are informational and the team doesn't process messages in them:

Section (topic)	What it's for
CEO	Write here. Everything you write goes straight to the team coordinator.
Tasks	Automatic task notifications — read-only.
Approvals	Requests for your approval — read-only.
System	System messages — read-only.

How to assign a task with /task

In the **CEO** section, type the `/task` command followed by the assignment:

```
/task Prepare 3 LinkedIn post variants about our new product XYZ
```

The task is created automatically, the CEO analyzes it and delegates it to the right specialist. You then follow the progress in the **Tasks** section of the web app — see the Tasks chapter.

How to just talk to the CEO

You don't need any command for a normal conversation. Simply write a question or request into the **CEO** section as you would to a colleague: "How is the team doing this week?", "I need materials for the investor meeting." The CEO replies right in Telegram.

Voice messages

Don't feel like typing? Record a regular Telegram voice message. The system transcribes it to text and the agent acts on it right away — just as if you had typed the request.

> 💡 **Tip:** Voice messages are great on the go: on your way back from a meeting, dictate a summary and tasks to the CEO, and by the time you reach the office, the team is already working.

Direct messages to the bot and user approval

You can also message the bot directly (outside the group). Who may message the bot is controlled by the **Direct Messages** setting in the Telegram administration:

- **Approved only** — the default and recommended mode. Only users approved by the administrator may write. A new person requests access first (see below).
- **Open (anyone can write)** — the bot replies to anyone who finds it.
- **Closed (DMs disabled)** — direct messages are off; only the group works.

Similarly, the **Group Chats** setting controls whose messages the team processes in the group: **Everyone**, **No one**, or **Approved only**.

Approving a new user works like this:

1. The new user (say, a colleague) sends the bot their first message. The bot replies that access isn't approved yet and the administrator's consent is needed.
2. The administrator opens **Settings → Integrations → Telegram**. In the **Approval Requests** section they see **Pending Requests** with the requester's name.
3. The administrator clicks **Approve** — the user gets a confirmation in Telegram and can communicate with the team from that moment on.

The **Allowed Users** list shows who has access; the administrator can revoke it at any time with the **Revoke** button.

✅ **Best practice:** Keep direct messages in the **Approved only** mode. You can be sure that only people you've approved talk to your AI team (and your company data).

How to connect Telegram (Administrator role)

1. In Telegram, create your own bot via the **@BotFather** bot and copy its token.
2. In Control Center, open **Settings → Integrations**, select **Telegram**, and click **Set up Telegram bot**.
3. Paste the token into the **Bot Token (from @BotFather)** field and click **Connect**.
4. Verify it works with the **Test Connection** button.
5. If you use a team group, fill in the **Group Chat ID** field and the numbers of the individual sections (**Coordinator Topic**, **Approvals Topic**, **Tasks Topic**, **System Topic**) in the **Routing** section, then click **Save Settings**.

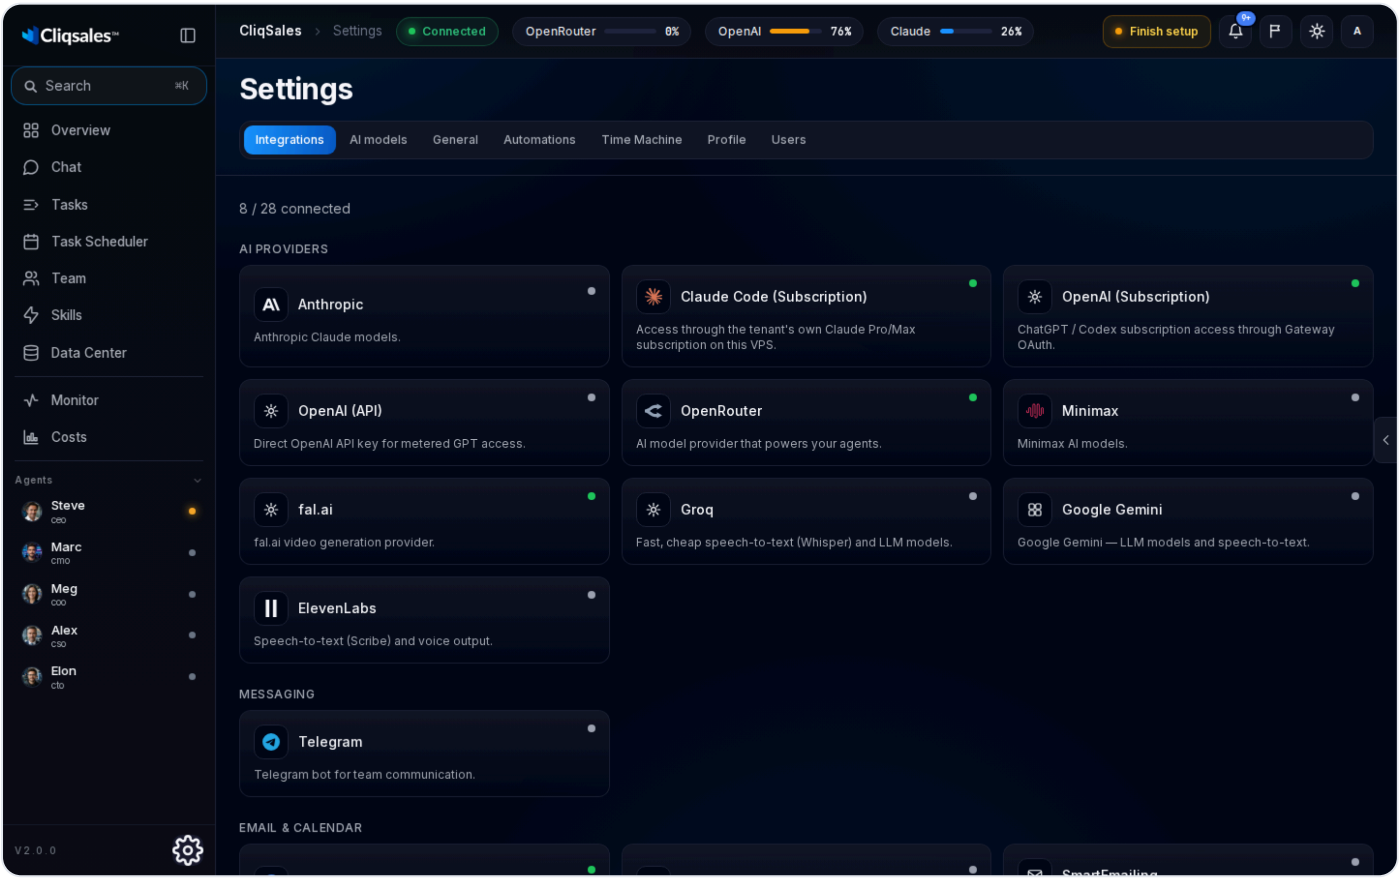

⚠️ **Watch out:** Connecting or disconnecting Telegram restarts the agents' runtime and may interrupt their work in progress. After saving, the app asks whether to restart now or later — the changes take effect only after the restart.

11 Integrations

Integrations connect your AI team to the tools you already use — your CRM, email, website, calendar. Once connected, agents work with them on their own: find a contact, schedule a meeting, publish an article. Everything is configured in **Settings → Integrations** (Administrator role); most connections are a one-time setup.

Finding your way around the Integrations section

Integrations are grouped into categories: **CRM**, **Messaging**, **Email & Calendar**, **AI Providers**, **Social & Marketing**, **Productivity**, **Operations**, **Search**, **Custom**, **Security**, and **Infrastructure**.

Each card shows the connection status:

- **Connected** — everything is filled in, the integration works.
- **Partial** — some details are missing; the integration may work only partially.
- **Not configured** — the integration isn't connected yet.

Click a card to fill in the credentials, save them with **Save**, and verify with the **Test connection** button (**Test Connection** on some integrations).

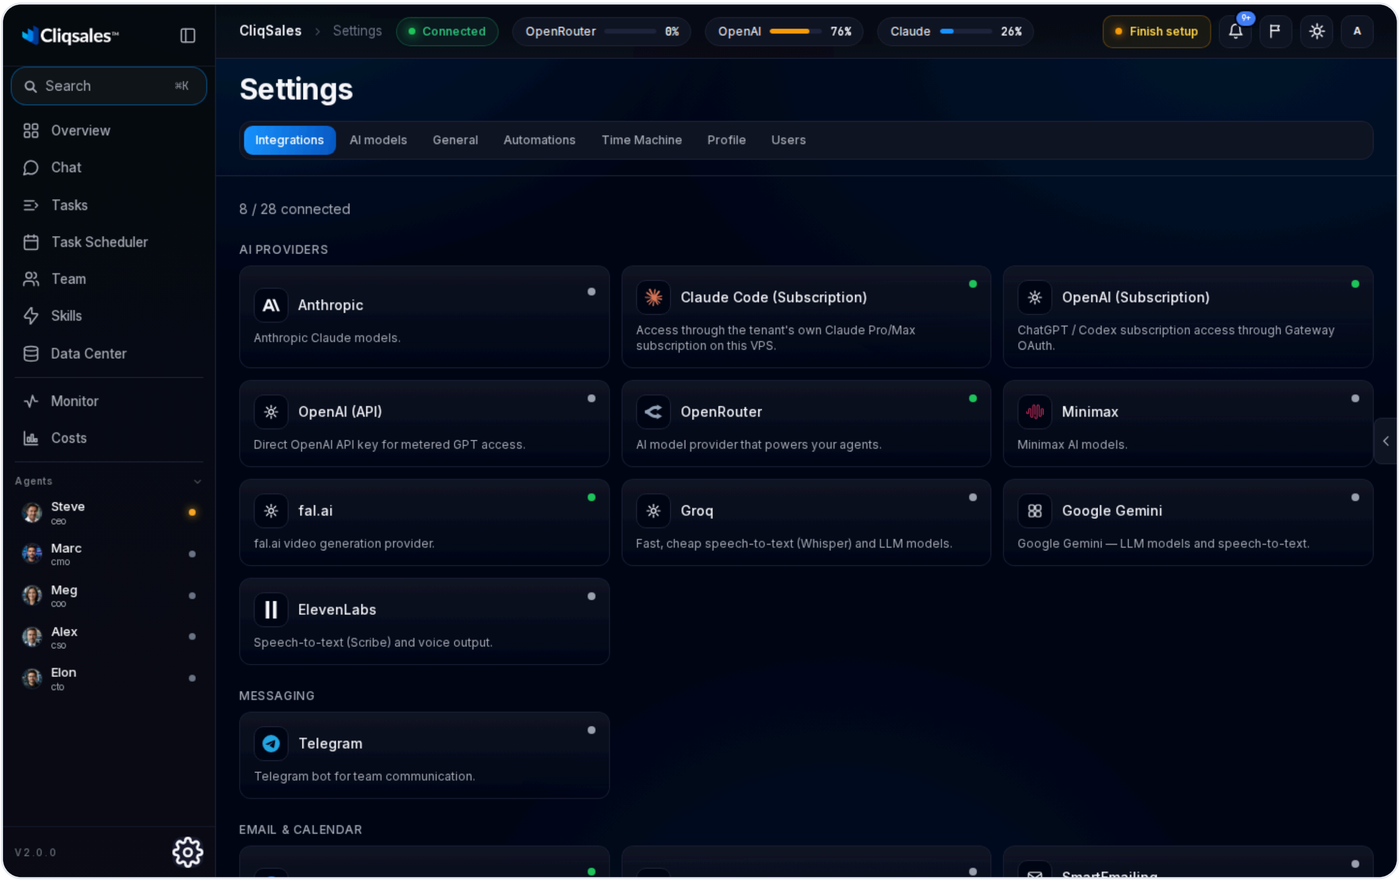

> ✅ **Best practice:** Give agents only the permissions they really need — with services where you choose an access scope (e.g. GoHighLevel), tick only the necessary items. And remember: you can revoke any token or key in the source service at any time, which ends the access immediately.

CRM — GoHighLevel

The **CliqSales (GoHighLevel)** card in the **CRM** category connects agents to your GoHighLevel account. Agents can then handle:

- **Contacts** — create, update, and look up clients,
- **Pipeline** — manage deals and move them through stages,
- **Calendar** — book meetings and watch availability,
- **Emails** — send emails through GHL,
- **Conversations** — read and reply to messages in the GHL inbox,
- **Social Planner** — schedule and publish posts to Facebook, Instagram, LinkedIn, Google Business Profile, and TikTok.

Then all it takes is: "Find all leads from last week and prepare a follow-up" or "Schedule a Facebook post for tomorrow at 10:00."

How to connect GoHighLevel (Administrator role)

1. In GoHighLevel, open **Settings → Integrations → Private Integrations** and click **Create New**. Name it e.g. "AI Team".
2. Select the permissions (scopes) per the table below, save, and copy the generated token (the Private Integration Token, PIT).
3. In Control Center, open **Settings → Integrations → CliqSales (GoHighLevel)**. Paste the token into the **Private Integration Token (PIT)** field and fill in the **Location ID** (in GHL: Settings → Business Info → Company ID).
4. Fill in the **Social media user email** — the GHL user under which Social Planner post drafts should be created.
5. Click **Test connection** — on success you'll see "Connected to location".

Permission (scope)	Purpose
`contacts.readonly`	read contacts
`contacts.write`	create and update contacts
`opportunities.readonly`	read deals
`opportunities.write`	move deals forward
`conversations.readonly`	read messages
`conversations.write`	send messages
`conversations/message.readonly`	read individual messages
`conversations/message.write`	send messages
`calendars.readonly`	read the calendar
`calendars.write`	book meetings
`calendars/events.readonly`	read events
`calendars/events.write`	create events
`campaigns.readonly`	read campaigns
`socialplanner/post.readonly`	read social posts
`socialplanner/post.write`	publish to social media
`locations.readonly`	read the location

> ✅ **Best practice:** Don't want agents changing deals? Don't give the token `opportunities.write`. Agents only ever have the permissions you grant the token, and you can revoke the token in GHL at any time.

WordPress

Connecting WordPress lets agents write and publish articles on your blog — including SEO optimization. How to assign articles and what the agent can do is covered in the chapter Content creation and media.

How to connect WordPress (Administrator role)

1. Sign in to your WordPress admin (yoursite.com/wp-admin), open **Users → Profile**, and in the **Application Passwords** section create a new password named "AI Team".
2. Copy the generated password right away — WordPress shows it only once.
3. In Control Center, open **Settings → Integrations → WordPress** and fill in the **Site URL** (e.g. `https://yoursite.com`, without a trailing slash), **User** (your WordPress username), and the WordPress **Application password**.
4. Save and click **Test connection**.

> 💡 **Tip:** Running several sites? The integration supports multiple profiles — in the **Saved profiles** section, add another site with the **New profile** button and mark one with **Use as default profile**. The profile's short **Identifier** is then used to pick the site when assigning work.

Google Workspace

Connecting Google gives agents access to selected services: Gmail, Calendar, Drive, Docs, Sheets, Slides, and optionally Contacts. It uses the standard Google sign-in (OAuth) — you never enter your Google password into Control Center.

How to connect Google Workspace (Administrator role)

1. Open **Settings → Integrations → Google Workspace**.
2. In the **Google account or Workspace email** field, enter the account you want to connect.
3. In the **Allowed services** section, choose which services agents may work with.

4. Click **Connect Google Workspace** — a new window walks you through the Google sign-in and consent.
5. After returning, verify the status with the **Test Connection** button.

> 💡 **Tip:** If the browser blocked the sign-in window, use the **Open Google sign-in** link and go through it manually.

Trusted-agent mode

By connecting, you allow agents to perform ordinary work actions in the allowed services without confirming every single step — which is why it's called "trusted-agent" mode. Dangerous or bulk operations are guarded by the tool safety layer.

Disconnecting

In the integration detail, click **Disconnect**. Control Center deletes the locally stored access and attempts to revoke it on Google's side as well.

When something doesn't work

Problem	Solution
The connect button is inactive	The installation doesn't have central Google sign-in completed — contact your installation administrator.
The Google sign-in didn't open	The browser blocked the new window. Use the **Open Google sign-in** link.
Test Connection failed	The access can't be refreshed. Disconnect the integration and connect it again.

Email accounts (IMAP/SMTP)

The **Email (IMAP/SMTP)** card in the **Email & Calendar** category lets you connect one or more mailboxes (e.g. info@, sales@, support@) as separate profiles. Agents can then read the latest messages from a mailbox on demand (only the inbox folder is read; there's no automatic background fetching) and — if you enable sending — send emails through that mailbox.

1. Click **Add mailbox**.
2. Fill in the **Display name**, **Email address**, **Purpose** (General, Sales, Support, Operations, Marketing, Technical), and **Owner**.

3. In the **Incoming IMAP** section, enter the **IMAP host**, **IMAP port**, **Security**, **Username**, and **Password / app password**.
4. Optionally turn on **Enable outgoing SMTP** and fill in the SMTP details, including the **From email** and **From name**.
5. Save with **Save mailbox** and verify with the **Test IMAP** and **Test SMTP** buttons.

The **Available to agents** switch controls whether agents may work with the mailbox.

> 💡 **Tip:** For Gmail and other services with two-factor authentication, use an app password generated in the service's account settings instead of your regular password.

Fireflies (meeting transcripts)

Fireflies.ai is a tool that records and transcribes your meetings. Once connected, agents have access to transcripts, summaries, and action items from meetings. Just fill in the API key on the card, save, and test the connection.

Meeting transcripts can also be stored in the company wiki automatically — every meeting becomes a wiki page. Setting up that flow (the inbound address for Fireflies and enabling processing in **Wiki → Settings**) is described in the chapter Team wiki.

SmartEmailing and Meta Ads

SmartEmailing (the **Email & Calendar** category) gives agents access to email campaigns and saved emails through the SmartEmailing API. Fill in the **Username** (usually your email) and **API key**, save, and test the connection.

Meta Ads (the **Social & Marketing** category) gives agents access to Facebook and Instagram ads. You need three values from Meta: a long-lived **Access token** (with the `ads_read` and `ads_management` scopes), the **Ad account ID** of your ad account (in the `act_...` form), and the **Page ID** of your Facebook page. Save and test the connection — a successful test prints the connected account and page.

fal.ai, HeyGen, and ElevenLabs

These cards in the **AI Providers** category extend the team's creative abilities — each just needs an API key from the source service:

- **fal.ai** — video generation (paid video models). You get the key after registering at fal.ai.
- **HeyGen** — music and sound effects for the video skills. Optional — without a key, videos are produced without music.
- **ElevenLabs** — speech transcription and voice output (e.g. voiceovers for videos).

What agents use them for and how to work with the results is covered in the chapter Content creation and media.

AI providers

For agents to think and reply, your team needs access to AI models. That's what the **AI Providers** category provides — the administrator sets it up, typically when the system is provisioned:

- **Claude Code (Subscription)** — the team uses your Claude Pro/Max subscription. The administrator pairs it with the **Connect Claude** button and a browser sign-in.
- **OpenAI (Subscription)** — a ChatGPT/Codex subscription is paired similarly (the **Pair ChatGPT/Codex** button, sign-in via a device code).
- **OpenRouter** — a credit-based add-on: you top up credit and the team draws from it based on usage. Works as a complement to a subscription or on its own.
- Alternatively, direct API keys can be entered (the **Anthropic**, **OpenAI (API)**, and other cards).

If no provider is connected, the app shows the warning "No model provider access configured" and agents can't work. What to do is covered in the Troubleshooting chapter; what operation costs and how to keep an eye on usage is in the Costs and limits chapter.

Automations: inbound webhooks

In **Settings → Automations** you connect the AI team to the outside world in the opposite direction: an external system (a web form, a booking system, another application) sends an event to your Control Center address — and that automatically triggers an agent's work. Such an inbound address is called a webhook: simply an address where another service sends a message.

Example: a new lead fills in a form on your website → the form sends an event → the agent, following your instructions, prepares a reply and creates a contact in the CRM.

How to create an automation (Administrator role)

1. Open **Settings → Automations** and click **Create webhook**.
2. Fill in the **Name**, pick an agent in the **Agent** field, and in the **Agent instructions** field write what to do with the incoming event. Incoming data cannot change these settings — the agent always follows your instructions.
3. Optionally pick **Skills the agent may use** and the **Context** — whether each event should be handled separately (**Each event separately**), continue per contact (**Same contact or deal together**), or all in one conversation (**One shared conversation**).
4. After saving, you get a **Public URL** and a **Bearer token** — enter both in the external system.
5. Verify it works with the **Send test event** button.

You can pause an automation at any time with the **Accept new events** switch.

> ⚠️ **Watch out:** The Bearer token is shown only once, right at creation — save it. You can replace it later at any time with the **Rotate token** button (the old one stops working).

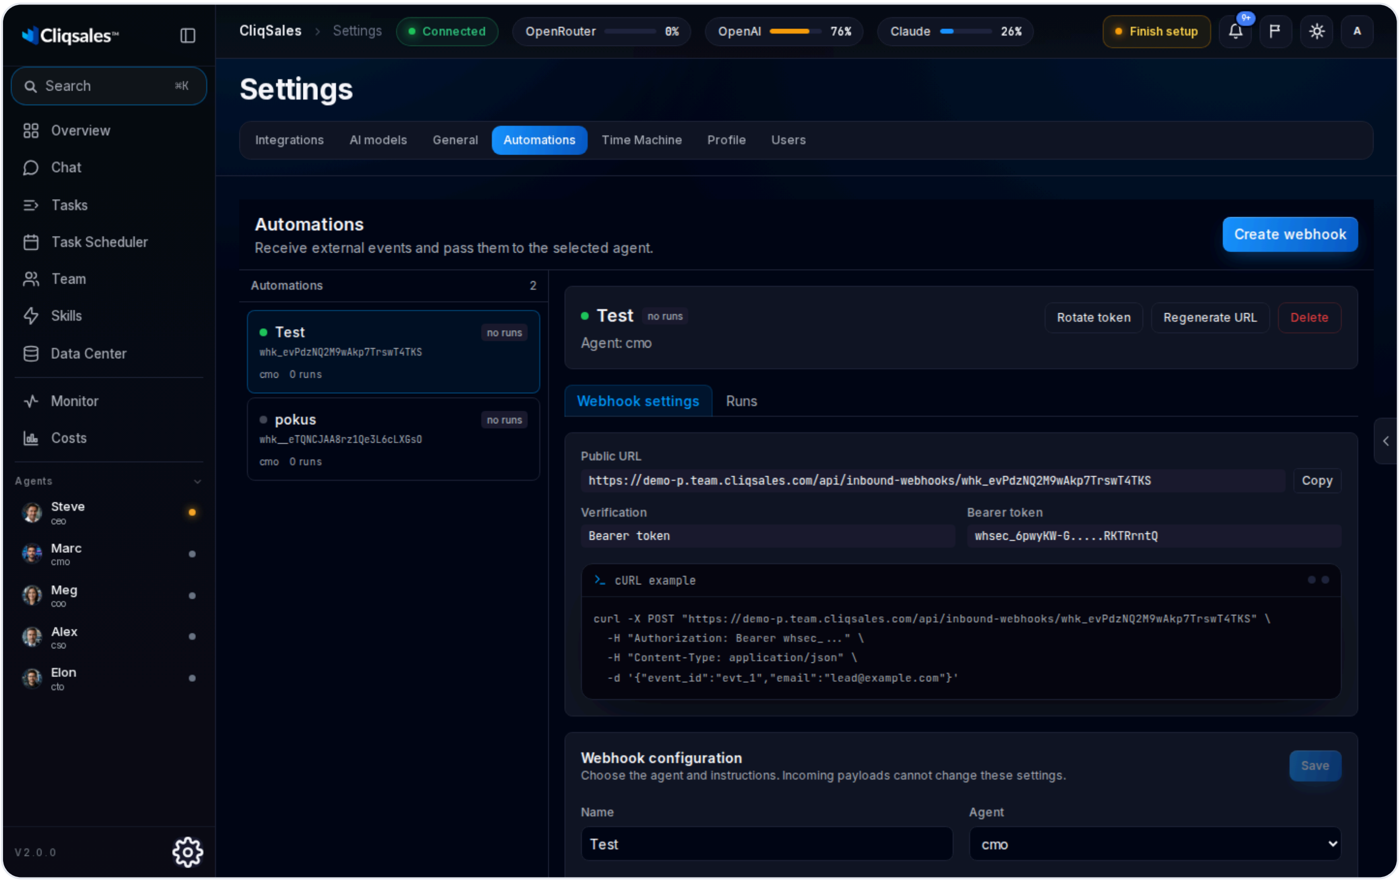

Run overview

The **Runs** tab shows every automation run with its status: accepted, queued, running, completed, failed, needs review, and more. A run's detail shows the timeline, the **Agent response**, and any **Outputs**.

Runs in the "needs review" state wait for your decision: you can close them manually or run them again. Running again (**Run again as new**) requires confirming that you've verified it's safe to repeat the event — for example that an email won't be sent to a customer twice.

Webhook security

Every webhook is protected by verification: a **Bearer token** (the external system sends a secret token) or an **HMAC signing secret** (the external system signs the request). Verification can also be turned off (**No verification**), but only for sources that can't send a token or a signature — the app asks for explicit confirmation that you understand the risk.

> ⚠️ **Watch out:** Without verification, anyone who knows the public URL can trigger the automation. If you had to turn verification off, treat the URL as a secret — and if you suspect it leaked, replace it with the **Regenerate URL** button (the old one stops working immediately).

12 Content creation and media

Your AI team can produce finished content: images, short videos, narrated text, blog articles, and social media posts. This chapter shows what you can ask the team for and how to phrase the request so the result is worth it. Connecting the required services (fal.ai, HeyGen, ElevenLabs, WordPress, GoHighLevel) is covered in the Integrations chapter — here we assume it's already done.

You assign everything the same way: write to an agent in chat, or create a task. Marketing content is the CMO's domain, but feel free to write to the CEO — they'll hand the work over themselves.

Images

Tell the agent what you want:

- "Create a banner for the spring sale Facebook campaign."
- "Prepare 3 hero image variants for the landing page."
- "Generate a product photo of a mug on a white background."

The agent generates the image and saves it among the team's files. It can handle photorealistic portraits and product photos, illustrations and icons, and marketing banners.

Generation runs through a configured model chain: the primary model is typically part of your AI subscription (no extra fees), and if it doesn't respond, the agent automatically tries the fallback models. You can fine-tune the model order, enabled fallbacks and the timeout in **Settings** → **Image generation** (Administrator role) — for everyday work there's nothing to set up.

You'll find the results in **Data Center** — the most convenient way to view new images and videos is the gallery view with thumbnails (the **Gallery** toggle). Details in the Data Center chapter.

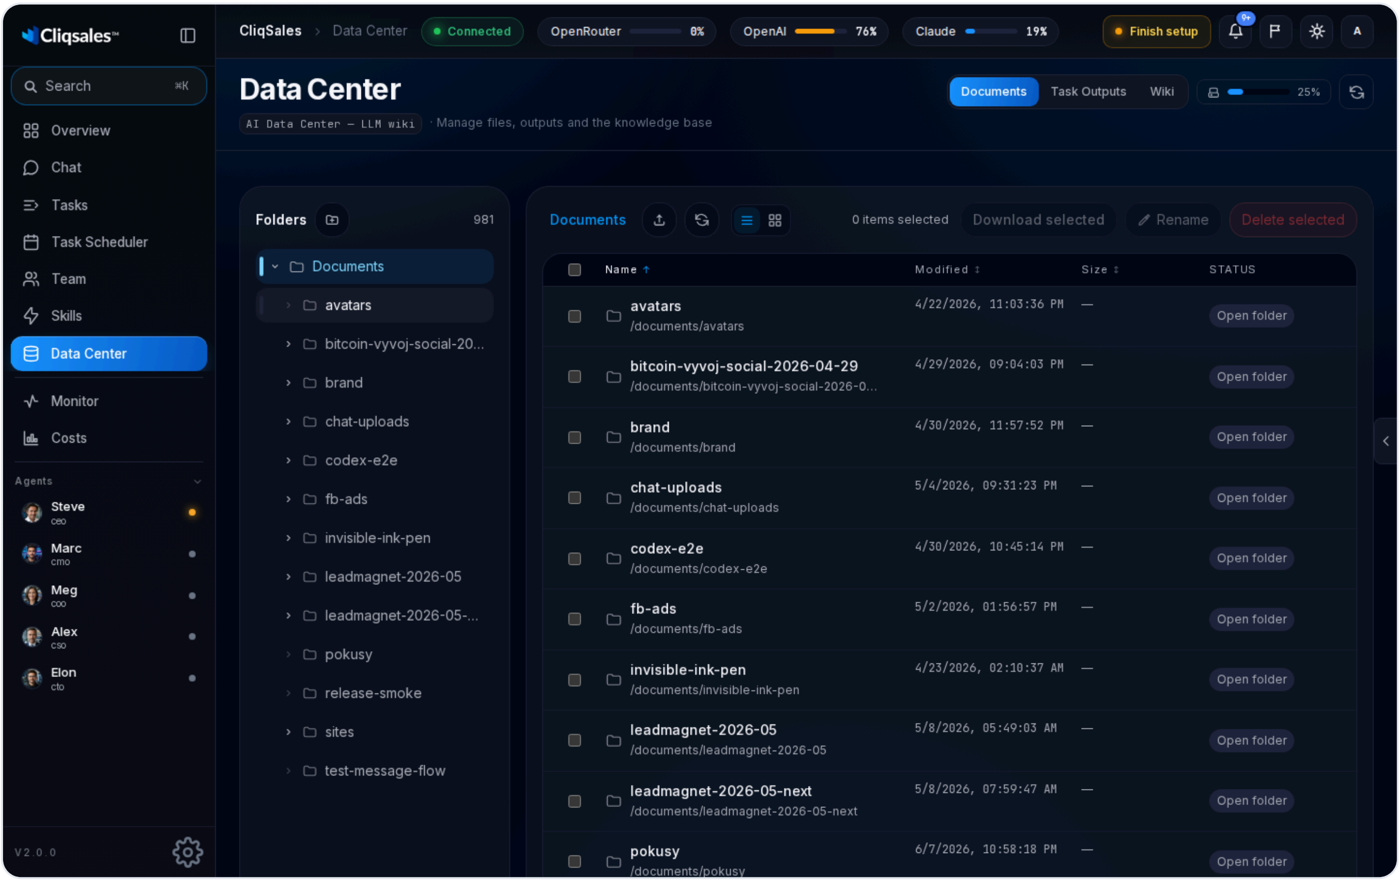

> 💡 **Tip:** Describe the style, not just the content: "photorealistic", "illustration", "minimalist", "in brand colors". The agent adapts. The more specific the description, the fewer revision rounds.
>
> ⚠️ **Watch out:** In older versions, the chat occasionally announced a finished image that was nowhere to be found. Since version 2.0, the chat announces completion only once the image is actually generated and saved.

Videos

The team has three video skills — depending on the kind of video you need. Just describe it in chat or in a task and the agent picks the right approach itself:

- **Live-action video** — real people on camera, including **Czech dialogue with lip sync**. The length isn't limited: a longer spot is built by chaining scenes with the same characters. Great for ads and spots with actors.
- **Animated video** — a viral animation (~20–60 s) in a style you pick: Pixar 3D, Ghibli/anime, claymation, LEGO, plush toy, or 2D motion. With a voiceover, or with talking characters with Czech lip sync.

- **Branded product video** — an elegant video built from your brand identity (30–90 s): 3D scenes, animated text, music. It can also do a "video from your website" — send a link to a page and the agent pulls its copy, photos, colors, and logo straight into the scenes. It renders locally, without paid video models.

How it works — and why it won't surprise you on cost:

1. The agent first prepares **concepts and a cheap storyboard with preview images** (it costs just cents) and waits for your approval.
2. Only after your approval does it start the **paid render** of the video. You approve every step — the more expensive generation never starts without your consent.
3. The finishing touches add subtitles (kinetic, word by word, never over a face), transitions, and music.

What you need for it: the **fal.ai** service connected (paid video models) and optionally **HeyGen** for music and sound effects — without a HeyGen key, videos are produced without music. Connect both following Integrations. Finished videos are in **Data Center** in the `videos` folder — the gallery shows only the final versions (`final-v1`, `final-v2`…); working files stay hidden.

The price depends on the length, engine, and number of scenes — thanks to approving the storyboard up front, you know what you're getting into, and you track the spend in the Costs and limits section.

> 💡 **Tip:** For videos, describe the shot like a director — camera movement ("slow zoom", "pan from left to right", "static shot"), light ("golden hour", "studio light", "neon glow"), and mood ("cinematic", "energetic", "professional"). A precise description = a better result.
>
> ✅ **Best practice:** The storyboard is your cheapest insurance. Give it a minute of attention and don't hesitate to have it redone — fixing an image costs cents, re-rendering a finished video costs dollars.

Voice and voiceover

Through the connected **ElevenLabs** service (see Integrations), the agent can turn text into a realistic voice — a voiceover for a video, an audio version of an article, or a voice message. Just say: "Record a voiceover for this video" or "Read this text aloud".

There are dozens of ready-made voices to choose from — male, female, various accents. On paid ElevenLabs plans you can even clone your own voice.

Approximate prices: ~0.07 Kč per sentence (100 characters), ~0.70 Kč per paragraph (1,000 characters), ~3.50 Kč per page (5,000 characters). The free ElevenLabs tier offers 10,000 characters per month at no cost.

Your content team in one sentence

You steer content creation with three sentences — the team takes care of everything else on its own and **never publishes without your explicit approval**:

1. **"Write me an article about X."** You get a complete concept: an article with SEO metadata, images including a preview card for social sharing, and a draft directly in your editorial system (WordPress or the CliqSales blog). You receive a link — review it and, if you like, say "publish".
2. **"Plan my content for 2 weeks."** The team looks into what your customers care about and presents a calendar with topics and reasoning. Once approved, it produces and schedules the drafts. A detailed plan makes sense for 1–4 weeks; for longer periods you get a thematic map, with the detail always covering the upcoming weeks.
3. **"Turn on daily content."** The team regularly prepares proposals on its own, following your rhythm, and sends them to you for approval — you just reply "OK". You can see and manage the recurring jobs in the **Task Scheduler** (pause, change the rhythm, turn off). "Pause content" works at any time.

For social media posts the team also produces finished visuals in your brand style — carousels (up to 8 slides), quote cards, or story covers, with the text right in the image. And when it writes a data-driven article and the source material is thin, it runs deep research and backs the numbers with cited sources.

> 💡 **Tip:** The system checks for completeness — a task with an agreed deliverable can't be handed in for review unfinished. If something is missing, the team gets a precise list and fills it in itself; at most you'll notice the result arrives a few minutes later, but complete.
>
> ✅ **Best practice:** Set up approvals to match your rhythm — one batch a week (the default), or a daily morning OK if you publish a lot.

Blog articles (WordPress)

With WordPress connected (see Integrations), the agent writes the article, optimizes it for search engines, and saves it to your site:

1. You say: "Write an article about the benefits of AI for small businesses."
2. The agent creates a structured article including SEO (search title, meta description, keywords), can assign categories and tags, and can upload the article's featured image.
3. The article is saved as a **draft** — it isn't published. You read it in WordPress, edit it, and publish it yourself.

If you want it published right away, say so explicitly: "Write **and publish** an article about..." If you have the Yoast SEO plugin installed, the agent sets the SEO title and meta description through it as well.

> ✅ **Best practice:** Articles under 300 words carry no weight with search engines — the recommended length is 800–1,500 words. For SEO, add the target keyword to the request: "optimize for the keyword: AI accounting automation".
>
> 💡 **Tip:** The agent knows your brand voice (your company's tone of communication) if you filled it in during the initial setup. Even so, it helps to restate the tone right in the request: "write matter-of-factly, no superlatives, address readers informally".
>
> ⚠️ **Watch out:** A request like "write something about us" is a reliable recipe for unusable text. Compare: "Write an 800-word article on 5 ways AI saves time in accounting, for small business owners" — a specific topic, length, angle, and audience.

Social media (GoHighLevel Social Planner)

Publishing to social media runs through the GoHighLevel Social Planner (for the connection, see Integrations). Supported platforms: **Facebook**, **Instagram**, **LinkedIn**, **Google Business Profile**, and **TikTok**.

Platform	Text	Image	Video
Facebook	yes	yes	yes
Instagram	yes	yes	yes
LinkedIn	yes	yes	no
Google Business	yes	yes	no
TikTok	yes	no	yes

Typical requests:

- "Prepare a post about our new service and schedule it for tomorrow at 10:00 on Facebook and Instagram."
- "Create a series of 5 posts for this week about AI transformation."
- "Schedule a LinkedIn post — the Novák client case study."

The agent prepares text optimized for each platform, adds hashtags and a call to action, and schedules the publication time. **You always have the last word:** the agent tells you about the prepared post, and you can approve it, edit it, or cancel it in the GHL Social Planner at any time. Without your approval, the agent doesn't publish.

> 💡 **Tip:** Instead of individual posts, assign series: "prepare a content calendar for next week" — the agent schedules posts for the whole week ahead and you approve them all at once. One post can go to several networks at the same time; the agent adapts the format to each platform.
>
> 💡 **Tip:** Give the agent context: the topic, target audience, and tone. "A post for e-shop owners, slightly informal, with a specific number from our case study" turns out better than "write a post about AI".
>
> ⚠️ **Watch out:** Never let posts go out without a fact check. The agent can phrase things convincingly even where it guessed a detail — always verify numbers, client names, and promises ("free shipping") before approving.

The agent can add images and videos to posts, but the media must be accessible via a URL (uploaded in GHL or another storage).

Where to find everything

Generated images, videos, and other files are saved by the team in **Data Center** — see Data Center. Articles are in WordPress (as drafts until you publish them) and scheduled posts in the GHL Social Planner.

13 Costs and limits

The **Costs** section shows how much work your AI team has done and what that work costs. It separates two layers: work running on a subscription and services paid as you go from credit. Use it when you want to know where the money goes, how much capacity is left, and whether limits are slowing the team down.

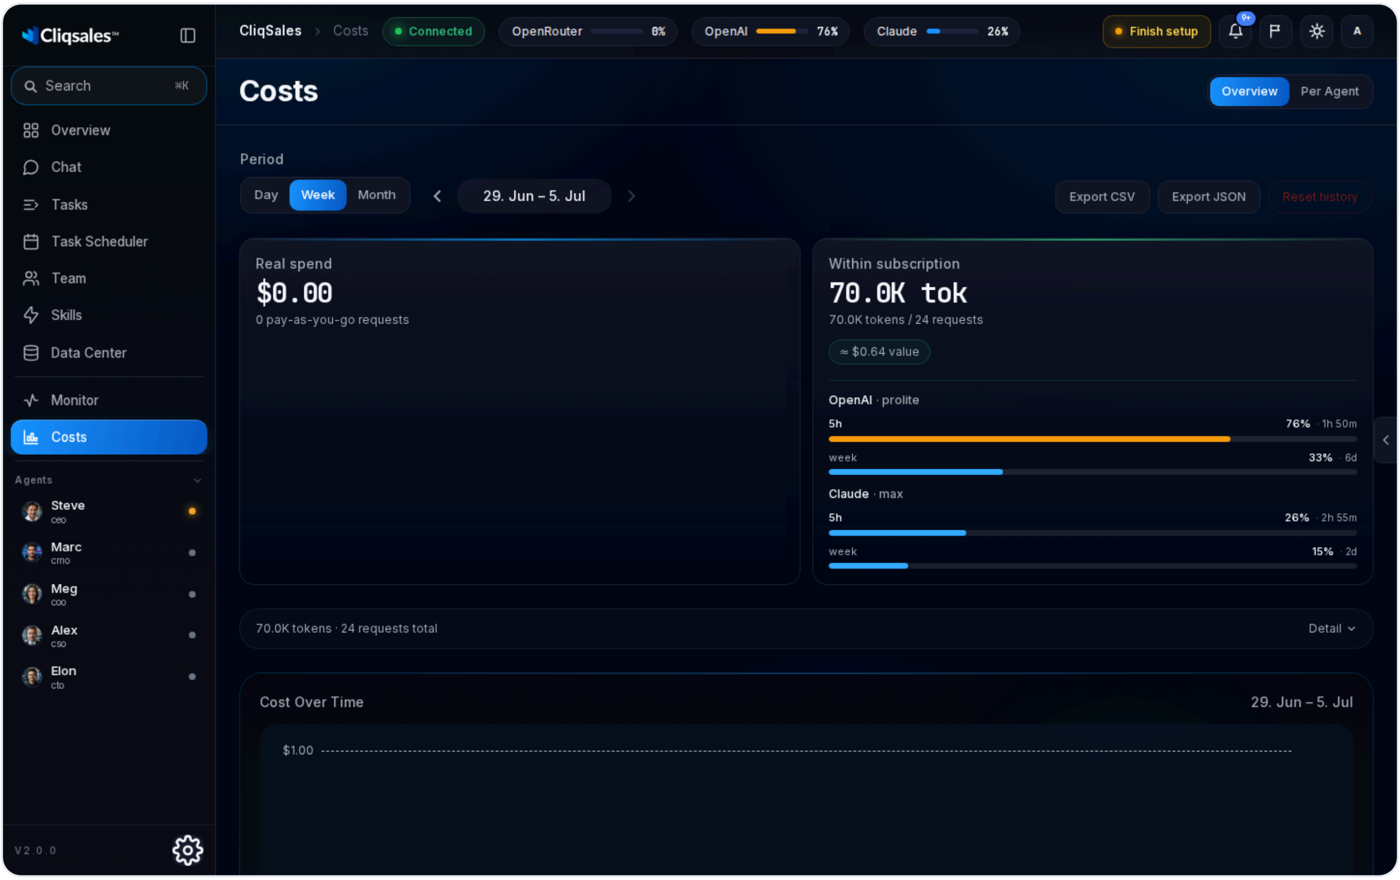

Two layers: subscription and pay-as-you-go

The panel distinguishes two billing modes that would otherwise blur into one number:

- **Within subscription** — work that runs on a prepaid AI plan. Individual requests aren't billed separately; they draw on the plan's capacity. The card shows the volume of work in tokens (a token is the unit language models use to measure the amount of text processed), the number of requests, and the approximate value of that work ("≈ … value").

- **Real spend** — what is actually deducted from credit or your wallet: typically OpenRouter, video generation, or voice. The card shows the amount for the selected period, the number of pay-as-you-go requests, and the breakdown by individual services.

So if your operation runs on a subscription, don't expect every request to show up as a new monetary cost — it draws on the plan's capacity, not credit.

Subscription usage gauges

A subscription has usage windows. On the **Within subscription** card you'll see two gauges for each connected plan:

- **5h** — what percentage of the 5-hour window is used up,
- **week** — the same for the weekly window.

Each gauge also shows how long until the window renews. The bar changes color: orange from 70% usage, red from 90% — a signal that the limit is approaching.

For OpenRouter (a credit service), you also see the spend and the limit reset time in the app's top bar — hover over it to see the detail **… left out of …** and **Resets in …**.

What happens when a window runs out

When a subscription window is exhausted, nothing is lost — work resumes after the limit resets, and you can see the reset time right next to the gauge. If the administrator has set up fallback models (**Settings → AI models**), the agent automatically switches to a fallback model in the meantime and keeps going. Since version 2.0, switching to a fallback model no longer comes with false error messages in the chat.

> 💡 **Tip:** Schedule big requests (bulk content generation, long research) away from the window's peak — ideally shortly after a reset, when the full capacity is free. You can see the reset time next to the **5h** and **week** gauges.
>
> ⚠️ **Watch out:** A common pitfall is launching several bulk jobs in parallel just before the window runs out. The window is exhausted mid-work and all the jobs then wait for the reset at once. Check the gauges before a big batch — above 70%, better to wait.

Service mix: where usage comes from

The AI team doesn't use just text. In the **By provider / service** breakdown, each item shows its service type: **Language**, **Images**, **Video**, **Voice**, **Transcription** (and possibly **Search** and **Other**). Items running on a subscription carry the **subscription** tag. That makes it easy to tell the agents' text work from multimedia — the latter is covered in the chapter Content creation and media.

The panel also offers:

- **Cost Over Time** — a chart of how usage evolves over the selected period,
- **Cost by Model** — which AI models carry the main load,
- the **Per Agent** tab — a cost breakdown by individual team members.

You switch periods with the **Day** / **Week** / **Month** buttons and the arrows between periods; you can download the data via **Export CSV** or **Export JSON**.

What billing paths your team has available

For agents to work, at least one model path must be active: either a paired subscription or an API key for a credit service (e.g. OpenRouter). Both are managed by the administrator in **Settings → Integrations** — see the Integrations chapter. If no path is active, the app announces it with a warning at the top and agents can't work.

Where the numbers come from

When a service provider returns the actual price directly in its response, Control Center uses it. When it doesn't, the price is calculated from the volume of work and a known price list, or estimated. So treat the panel as an **operational overview** of your usage, not as real-time accounting truth — always verify exact billing with the service provider.

How to read the panel in practice

"Where is our credit going?"

Look at the **Real spend** card and its breakdown by service, then at the **By provider / service** breakdown. The service type reveals whether credit is going to text, images, videos, or voice.

"Is the subscription slowing us down?"

Look at the **5h** and **week** gauges on the **Within subscription** card. If you're repeatedly getting close to 90%, the team is hitting its capacity ceiling — it helps to spread work over time, set up fallback models, or move to a higher subscription plan.

"Which agent uses the most?"

Switch to the **Per Agent** tab — you'll see the cost breakdown per agent, including models and the daily distribution.

If the numbers don't match your expectations or agents report a limit, see the Troubleshooting chapter.

14 Backups — Time Machine

Time Machine creates a backup of your platform's entire state with one click. When something goes wrong — you delete data by mistake, an agent makes an error, or you want to roll back changes — you restore the state from any point in the past. You'll find it in **Settings** (the gear at the bottom left) → **Time Machine**.

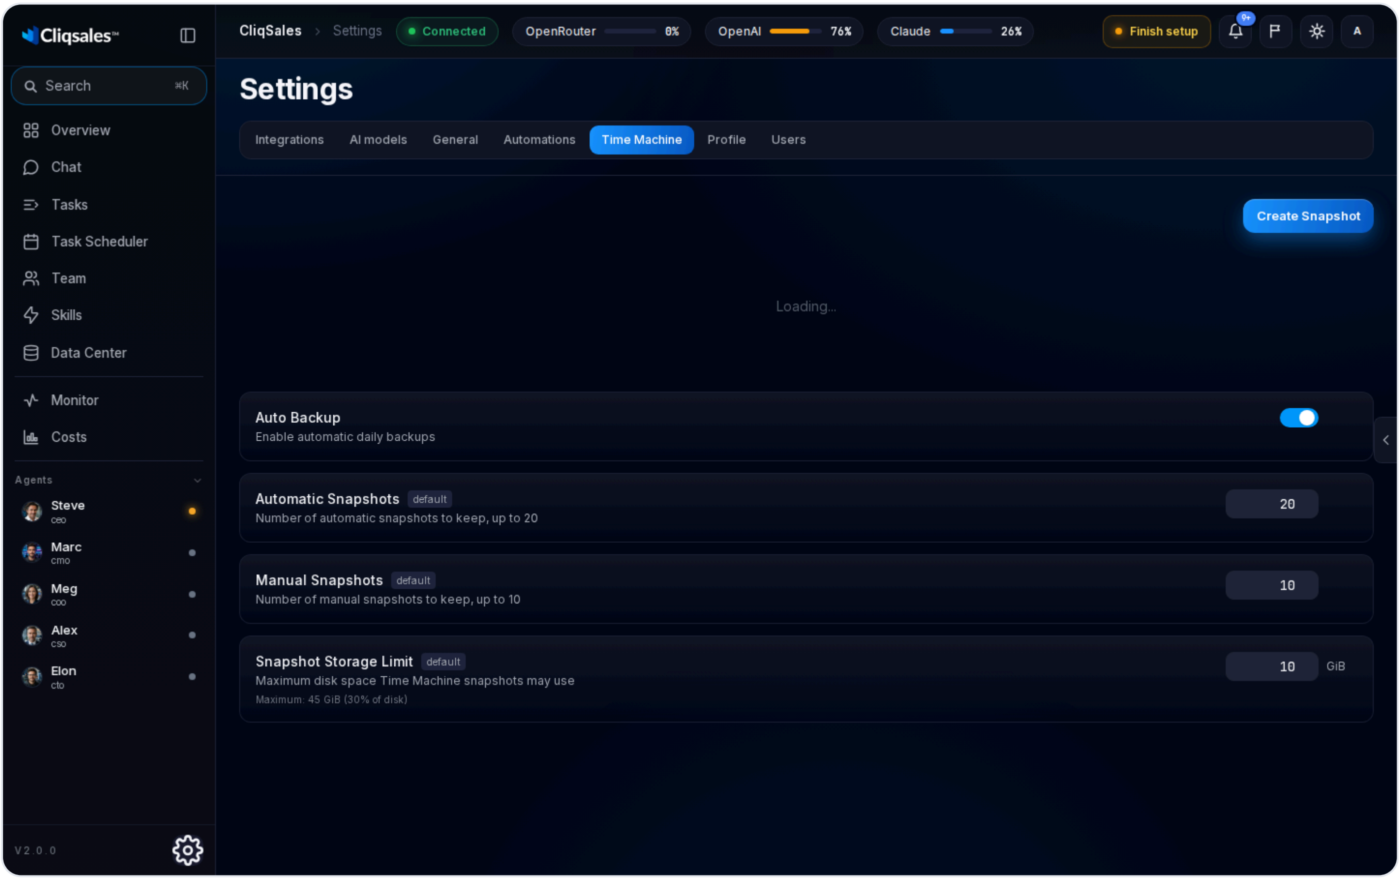

How to create a snapshot

1. Open **Settings → Time Machine**.
2. Click **Create Snapshot**.
3. Enter the **Snapshot name** — name it after the situation, e.g. "Before the contact import".
4. Confirm with **Create**. The snapshot is created on the server and appears in the list with the **Manual** badge.

> ✅ **Best practice:** Create a manual snapshot before every bigger change — a new workflow, a bulk import, settings changes, or a big document cleanup. A snapshot is free and takes a moment; getting back without one is much more expensive.

How to restore from a snapshot

1. In the list, find the snapshot you want to return to.
2. Click **Restore** and confirm. Before restoring, the current state is automatically backed up as a safety copy — so you can also come back "to now".
3. After the restore, the page reloads on its own after a moment; if it doesn't, refresh it in the browser manually. A restore typically takes a short while, depending on the size of your documents and the agents' state.

Only a user with the **Admin** role can perform a restore (see Administration and security).

> ⚠️ **Watch out:** A restore replaces the platform's current data with the snapshot's state — everything at once, not just a selected part. Everything created between the snapshot and now disappears from the regular view. Thanks to the automatic safety copy you can come back; however, only the last two safety copies are kept, so don't put off going back.

Automatic snapshots

You don't have to look after regular backups; the system creates them itself:

- **once a day** at night (around 3:00 server time) — they carry the **Auto** badge in the list,
- **before every restore** and before a reset — internal safety copies that don't show in the regular list.

How many snapshots are kept

Manual and automatic snapshots are deleted separately; older ones rotate automatically:

Type	How many are kept
Manual snapshots	configurable, at most 10
Automatic snapshots	configurable, at most 20
Internal safety copies	2 (not visible in the regular list)

You set the counts and the overall **Snapshot Storage Limit** right in **Settings → Time Machine**; the storage limit cannot exceed 30% of the server's disk. Snapshots cost nothing extra — they only take up space on the server.

Remove an unneeded snapshot with the **Delete** button.

What a snapshot contains

- all tasks and their state,
- settings and configuration,
- the chat history with agents,
- uploaded documents and the team wiki,
- user accounts and permissions,
- the agents' state (their working files, skills, and access).

For each snapshot in the list you can see which parts it covers (e.g. **Database**, **Documents**, **Wiki**, **OpenClaw** — the agents' state).

What a snapshot does not contain

- **CRM data** — contacts, deals, and meetings live in GoHighLevel, which has its own backups,
- external data outside your server (e.g. the content of connected third-party accounts).

Frequently asked questions

How often should I back up? Manually before every bigger change; for normal operation the daily automatic snapshots are enough, plus perhaps one manual snapshot a week.

Can I go back to a snapshot from a week ago? Yes, if it still exists — manual and automatic snapshots are deleted separately per the retention above.

Can I download a snapshot? Snapshots are stored on the server. For a manual download, contact your administrator.

What happens to the data between the snapshot and now? It will be replaced. But an automatic safety copy is created before the restore, so you can also return to the current state.

15 Administration and security

This chapter is mainly for whoever looks after Control Center at your company: how to add colleagues and control who may do what, how everyone manages their own password, and how the security of your data works.

Users and roles

Users are managed by the administrator in **Settings** (the gear at the bottom left) → **Users**. The section has two tabs: **Management** (accounts and access requests) and **Audit Trail** (a record of important events).

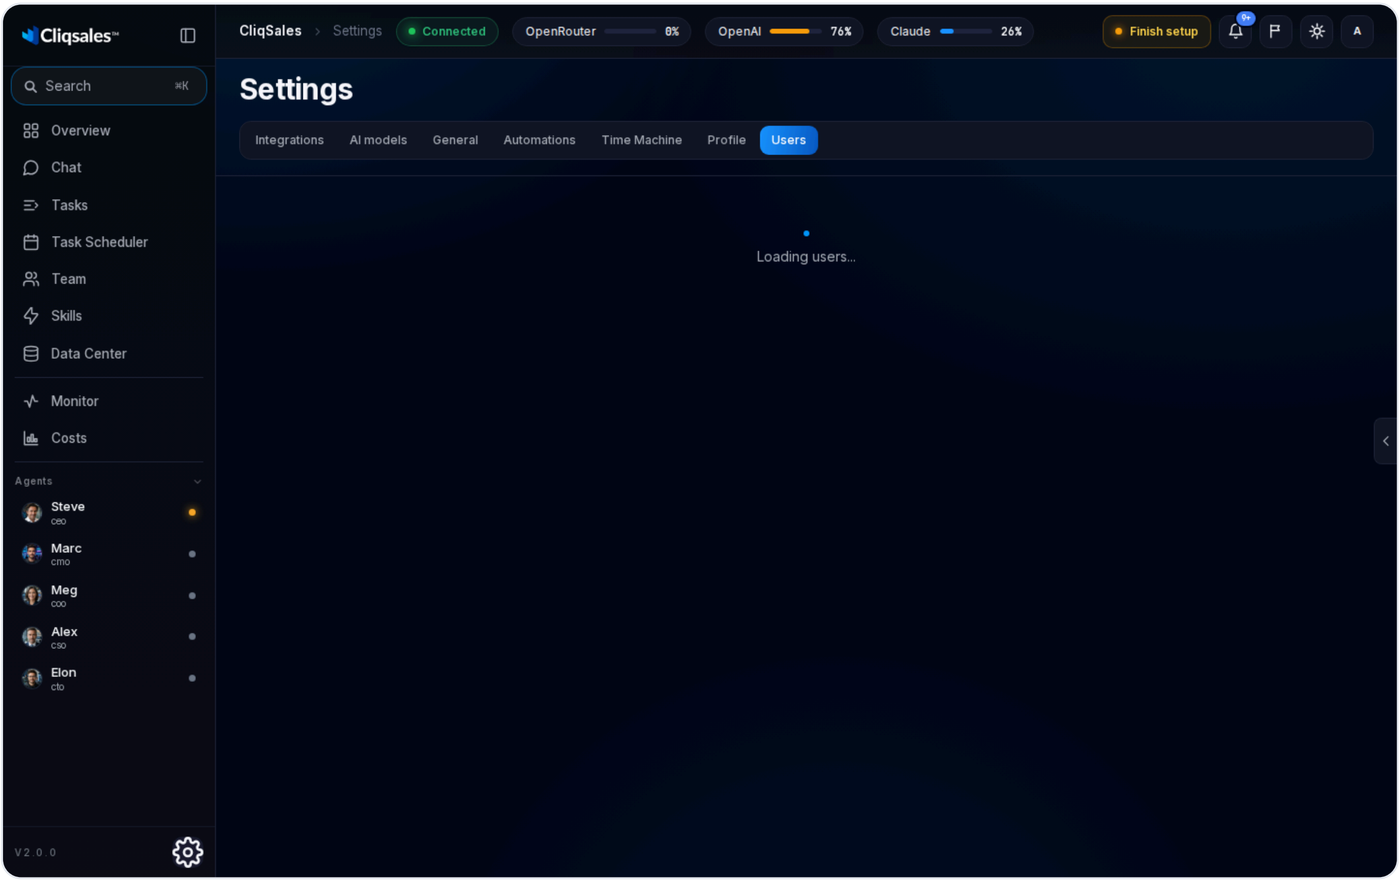

Each account has one of three roles:

Role	What it may do
Viewer	Look only: browses the board, tasks, conversations, and overviews. Can't change anything or assign work to the team.
Operator	Regular day-to-day work: chats with agents, creates and edits tasks, works with documents and the wiki, creates snapshots.
Admin	Everything an Operator can, plus administration: settings, integrations, user management, and restoring from a snapshot.

> ✅ **Best practice:** Assign roles by actual need — don't make everyone an administrator. A colleague who only reads the results is fine with **Viewer**; whoever works with the team daily needs **Operator**. Keep the number of administrators as low as possible: each of them can change the settings of the whole platform.

How to add a user

1. Open **Settings → Users**.
2. Click **+ Add Local User**.
3. Fill in the **Username**, **Password**, and **Display name**, and pick a role (**Viewer** / **Operator** / **Admin**).
4. Confirm with **Create User** and hand the sign-in details to your colleague. They can then change the password themselves (see below).

Editing and removing a user

In the user list, click **Edit** — you can change the display name and the role, and also set a new password (the **New password (optional)** field). An administrator can't change their own role, so they don't accidentally lock themselves out.

The **Delete** button removes an account — useful when a colleague leaves or no longer needs access.

Audit: who did what

The **Audit Trail** tab in the **Users** section records important events: sign-ins including failed attempts, sign-outs, password changes, user creation, edits, and deletion, snapshot creation and deletion. Records can be filtered by action type and by actor — handy when you need to find out after the fact who changed what and when.

Your profile and password

- **Changing your password:** click **Account** in the top-right corner and choose **Change password**. You enter your **Current password**, **New password**, and **Confirm new password** — the new password must be at least 8 characters long.
- **Forgotten password:** on the sign-in page, use the **Forgot password?** link — a reset link is sent to the email address on the account. If the email doesn't arrive, ask an administrator — they can set a new password for you in **Settings → Users** via **Edit**.
- **Display name:** an administrator changes it in **Settings → Users** via **Edit**.

> ✅ **Best practice:** Use strong passwords — 8 characters is the minimum, better is a longer phrase you don't use anywhere else.
>
> ⚠️ **Watch out:** Don't share one account between several people. The audit then won't show who really did what, a password change "kicks out" everyone at once, and permissions can't be told apart. Every colleague should have their own account with their own role — adding one takes a minute.

How the security of your data works

Without the technical details — this is how your platform is built:

- **Your own server.** Your Control Center instance runs on its own server, separated from other clients. Your data isn't shared with anyone.
- **Integrations via official APIs.** Connected services (CRM, email, social media…) use the official interfaces of the given platforms with your access tokens (a token is a key you use to grant a service access). You decide the permissions yourself and can revoke a token at any time.
- **Encrypted channels.** Communication channels (e.g. Telegram) use the platforms' own encryption.
- **Nothing goes out without you.** Outputs — emails, posts, articles — are prepared by the agent, but you approve the sending or publication. The only exception is when you explicitly order publication ("write and publish").
- **AI can hallucinate.** Artificial intelligence occasionally "makes up" a fact. That's why you always check the facts in important outputs — the final decision is yours.

You'll find an overview of connected services in **Settings → Integrations** — see the Integrations chapter.

> ✅ **Best practice:** Every now and then, go through **Settings → Integrations** and check that only the services you actually use are connected. Disconnect unused integrations — fewer accesses mean less risk.

What the AI team doesn't do

For realistic expectations:

What the AI team doesn't do	Why
Send emails on its own without your approval	Security — you always approve the outputs
Access the internet on its own initiative	It uses the internet only when it has an assignment and an enabled integration
Guarantee 100% factual accuracy	AI can hallucinate — check the facts
Replace human judgment	It's a tool; you make the final decisions

Related chapters: Backups — Time Machine for data protection, Troubleshooting for sign-in issues.

16 Troubleshooting

A first-aid chapter: the most common situations where something doesn't behave as expected, and what to check in what order. You'll solve most problems yourself within a minute — and when you don't, the end of the chapter tells you how to contact support and what to attach to a report.

> 💡 **Tip:** The fastest diagnosis is often asking the CEO in **Chat**: "What's happening with task X right now?" or "Why is nobody answering me?" Agents often know about the system's state and answer in plain language.

An agent in chat isn't responding

1. Look at the **agent's status**: in the **Team** section (or in the **Agents** list at the bottom of the left menu), each agent has a colored dot and a status — **Idle**, **Busy**, **Offline**, or **Error**. An agent that is **Busy** may reply with a delay; **Offline** or **Error** means a problem.
2. Check the **connection status** in the top bar. A green **Connected** is fine. Orange states (**Starting…**, **Connecting…**, **Restarting…**, **Updating…**) are transient — wait a moment. A red **Disconnected** means a connection outage.
3. Refresh the page in the browser (F5). Conversations aren't lost — threads survive a page refresh.
4. Look at the **work progress** in the conversation: for longer requests, the agent continuously shows which steps it's taking. Simple replies take seconds, bigger tasks minutes.
5. If an orange banner about missing model access shows under the top bar, agents can't work — see below.

> ⚠️ **Watch out:** Clicking the green connection status pill restarts the AI core (the **Restart OpenClaw** button). The restart interrupts all agents' work for a few dozen seconds — use it only as the last step when the team hasn't responded for a long time.

A menu item is gray with a "Requires gateway" note

The **Team** item in the left menu needs a live connection to the agents' runtime (the gateway). When that connection restarts or drops briefly, the item grays out and the tooltip shows **Requires gateway**.

1. Look at the **connection status** in the top bar — it most likely shows an orange transient state (**Starting…**, **Connecting…**).
2. Wait a moment. Once the connection is back (a green **Connected**), the item works again on its own.
3. If the connection doesn't recover for a long time, refresh the page in the browser.

Orange banner: model access is missing

The banner *"No model provider access configured. Agents cannot process messages. Connect a model provider in Settings → Integrations."* means that no model path is active — no subscription and no API key. Agents can't work until then.

1. In **Settings** → **Integrations** (requires the Administrator role), either pair a subscription (e.g. **Pair ChatGPT/Codex**, the Claude sign-in) or add and verify an **OpenRouter** key.
2. Once connected, the banner disappears and agents start processing messages.

The detailed procedure is in the Integrations chapter.

An agent reports a subscription limit

An AI model subscription has usage windows: **5-hour** and **weekly**. When a window is exhausted, work pauses and resumes after the limit resets.

1. Open the **Costs** section and look at the **Within subscription** card — the usage gauges for the 5-hour and weekly windows and the time when the limit renews.
2. The same gauges are also in the top bar (usage pills with a "resets in …" note).
3. If a fallback model is configured (e.g. via OpenRouter), the agent can continue on it in the meantime.
4. Wait for the window reset — work then resumes on its own; nothing is lost.

Details on windows and billing are in the Costs and limits chapter.

The Telegram bot doesn't react

1. Check that you're writing in the **right topic** — assignments belong in the CEO (coordinator) section. Messages in the other topics (Tasks, Approvals, System) are for notifications and approvals.
2. Check that the bot is a **member of the group** and has **administrator** rights — without them it doesn't see group messages.
3. In **Settings** → **Integrations** → Telegram, click **Test Connection** — it verifies the bot token and the connection.
4. Check the **user approval**: only approved users have access to the bot. You'll find them under **Allowed Users**; confirm pending **Approval Requests** with the **Approve** button.

Detailed setup is in the Communication channels chapter.

The password reset email didn't arrive

1. Check your **spam** folder (or "bulk" / "promotions").
2. Check that you entered the right value in the **Username or email** field — for security reasons the app shows the same confirmation even if no account exists for the entered address.
3. Keep in mind that the **link is valid for 2 hours** and can be used only once. An expired link is flagged by the page right away, offering **Request a new link**.
4. If the email doesn't arrive, an **administrator** can change your password: in **Settings** → **Users** they open your account and set a new password.

The full password recovery procedure is in the Getting started chapter.

Wiki: a page didn't appear or a source didn't load

The first step is always the same: open **Data Center** → **Wiki** → **Jobs** and find the source in question.

1. **A page didn't appear** — check the source's status. **Pending** means it's waiting its turn; for the **Failed** status, open the timeline and look at the last error.
2. **A Fireflies meeting didn't arrive** — check the Fireflies API key in Integrations, the inbound address (webhook) and the secret, and that Fireflies is enabled in **Wiki** → **Settings**.
3. **An email didn't load** — verify that the email account is added in Integrations, that the right folder or Gmail labels are selected, that the source is saved in the

wiki settings, and go through the result of the no-write test.
4. **The daily overview wasn't generated** — check that it's enabled and that the time and time zone are set; the overview history is in the wiki settings.

A detailed breakdown of all these situations is in the Team wiki chapter.

A file upload failed

The most common cause is size limits:

1. **Chat attachments:** at most **4 files** per message; images up to **5 MB**, documents up to **20 MB**, audio and video up to **100 MB**.
2. **Data Center → Documents:** documents up to **20 MB**, archives and media up to **100 MB**.
3. Check the **file type** — the error message lists the supported formats (PDF, DOCX, TXT, MD, common images, videos, and audio).
4. Check the **free space**: **Data Center** has a **Disk** indicator with the used space on the server. When it's filling up, delete unneeded files.
5. The "Upload was interrupted" message points to an internet connection outage — check your network and try again.

> 💡 **Tip:** A large file can often simply be made smaller — save an image at a lower resolution, trim a long video, send several documents in two messages.

A task hasn't moved for a long time

1. Open the task detail and the **Conversation** tab — the agent usually describes there what it's waiting for (your approval, a missing document, an answer to a question). Reply right in the conversation.
2. Check whether the task is sitting in the **Review** column — then it's waiting for your **Approve** / **Rework** decision.
3. Watch for the orange "Warning: N blocked tasks" banner above the board — the **Show** button opens just the blocked tasks. In the detail you'll find the **Block reason** box, and once the cause is removed you can get the task going again with the **Retry** button.
4. Ask the CEO in **Chat**: "What's blocking task X?"

More on task states in the Tasks chapter.

I need to roll back data after a mistake

When something gets deleted or broken by accident, use the backups:

1. Open **Settings** (the gear at the bottom left) → the **Time Machine** tab.
2. Pick a snapshot from before the problem and click **Restore**. Before restoring, the current state is automatically backed up, so you can also come back.
3. After the restore, just refresh the page in the browser.

Details in the Backups chapter.

I want to see what's happening right now

When you're not sure whether the team is working, you have two windows into the operation:

- **Monitor** in the left menu — the **Activity** (what's happening in the system), **Logs** (detailed records), and **Sessions** (the agents' running work blocks) tabs. Useful mainly for administrators and curious users.
- The **Live Feed** — a panel at the right edge of the screen with team events in real time. If the panel is closed, open it with the arrow button at the right edge (**Show live feed**).

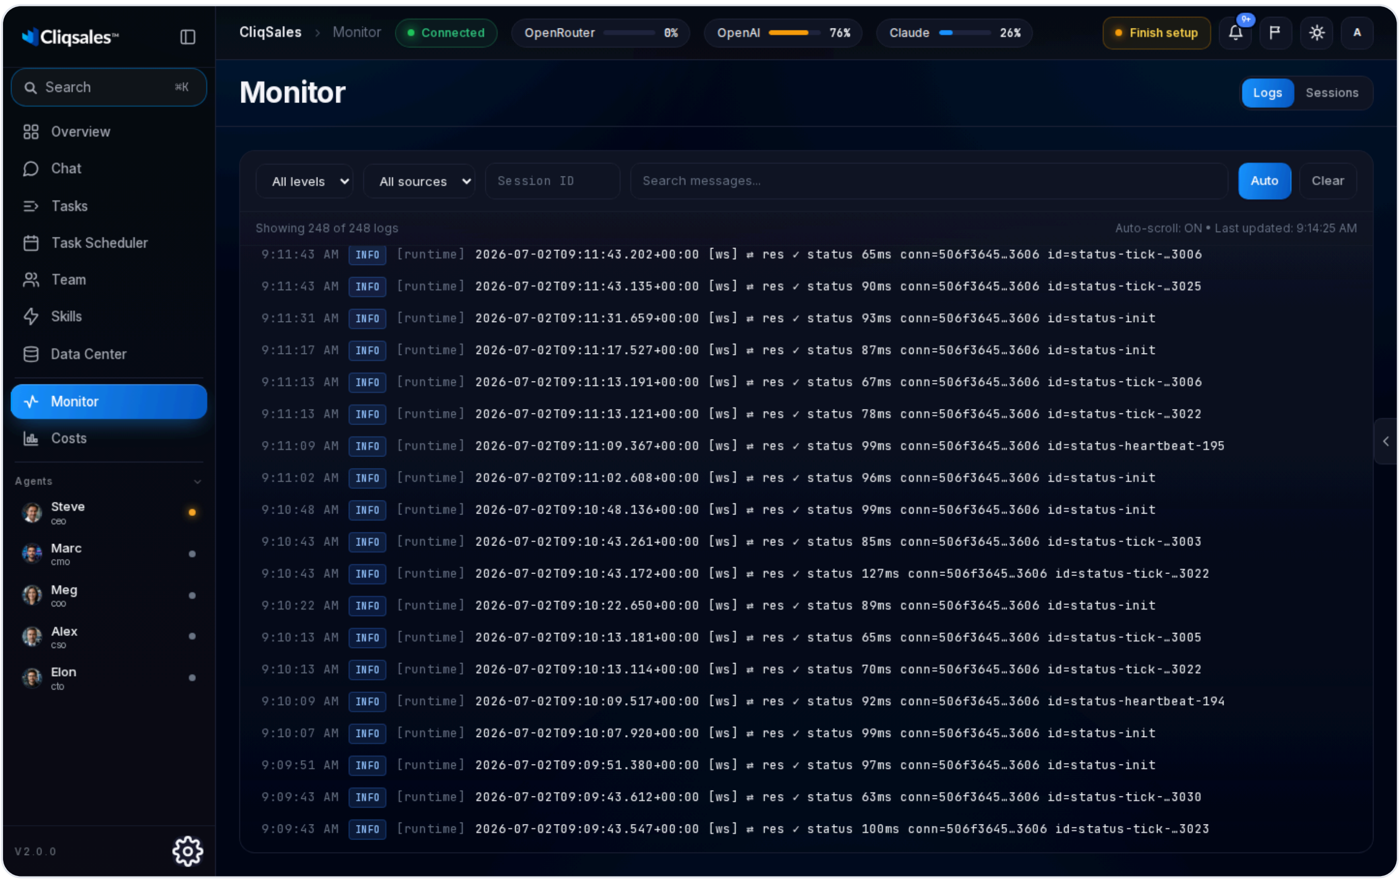

> ✅ **Best practice:** Before you restart anything, watch the **Live Feed** for a minute. If events keep coming in, the team is working — it's just taking time.

Contacting support

When you're stuck, get in touch:

- **Email:** support@cliqsales.cz
- **Telegram:** @krepaclawbot

So we can help quickly, please attach:

1. **A description of the problem** — what you expected and what happened instead.
2. **The steps** — what exactly you did (the section, buttons, the text of your request).
3. **A screenshot** of the screen with the problem, ideally including the error message.
4. **The time** the problem occurred (approximately is fine), plus the name of the task or agent involved, if any.

> 💡 **Tip:** You can report a bug quickly even without email: the **Feedback** icon in the top bar opens a window that sends the report straight from the app.

17 Glossary

A short alphabetical overview of the terms you'll come across in Control Center and in this handbook. Each term is explained in plain language, in one or two sentences; terms with their own chapter link to the details.

> 💡 **Tip:** When you hit an unfamiliar term right in your work, ask the CEO in **Chat** — they'll explain it in the context of your business.

Agent — an AI specialist on your team. It has its own role, memory, and skills; you write with it in **Chat** and it completes tasks. See Your AI team.

Agent roles (CEO, CMO, CSO, COO, CTO) — your team's five specializations: the CEO manages and delegates (your main contact), the CMO does marketing, the CSO sales, the COO operations and reports, the CTO technology and integrations. See Your AI team.

AI team — the five agents (CEO, CMO, CSO, COO, CTO) who work for you 24/7 — from marketing copy to reports. See Your AI team.

Board — the board in the **Tasks** section: columns by work status, with task cards you drag between them with the mouse. See Tasks.

Brand voice — your brand's voice: the tone and style your company communicates in (e.g. "professional but friendly"). You enter it in the setup wizard (**Tone of voice**, **Brand guidelines**) and agents stick to it in their copy.

Card (wiki) — one wiki page (a person, company, project…). In the **Graph** section each card is a circle; the lines between them show relationships. See Team wiki.

Channel (communication) — a way to talk to the team outside the app, for example Telegram. See Communication channels.

Credit vs. subscription — two ways of paying for AI services. Credit is deducted as you go per individual call (e.g. OpenRouter, videos, voice); a subscription is a flat plan drawn from within time windows. See Costs and limits.

Data Center — your team's storage with three tabs: **Documents** (your source materials), **Task Outputs** (what the team created), and **Wiki** (company memory). See Data Center.

Deliverable (output) — a task's concrete result: a file, text, link, or report. You'll find them in the task detail and collected in **Data Center → Task Outputs**. See Tasks.

Drag-and-drop — upload files by dragging them with the mouse straight into the chat window or a folder in **Data Center**; drag task cards between the board's columns.

Entity (wiki) — a concrete thing the wiki recognizes in sources: a person, company, project, product, process, or system. Every entity has its own page. See Team wiki.

fal.ai — an external service agents use to generate videos (from text and by animating an image). It needs an API key in **Settings → Integrations**; you pay with credit per generated video.

Fireflies — a service for automatic meeting transcripts. Once connected, transcripts can flow into the team wiki automatically. See Team wiki.

Gateway — the connection between the app and the agents' runtime. When it restarts, the **Team** item in the menu temporarily grays out with a **Requires gateway** note. See Troubleshooting.

GoHighLevel — a CRM platform that can be connected to Control Center. Agents then work with contacts, deals, the calendar, and social media post scheduling. See Integrations.

Hallucination — when AI "makes up" a fact that sounds credible but isn't true. That's why you always check the facts in important outputs — the final decision is yours.

Integration — a connection between Control Center and an external service (Telegram, WordPress, Google Workspace, CRM…). Managed in **Settings → Integrations**. See Integrations.

Live progress — watching the team work in real time: the agent's steps right in the conversation, the **Live Feed** panel at the right edge of the screen, and the **Monitor** section. See Finding your way around.

Model (AI) — a language model, the "brain" that thinks and writes for an agent. Access to models is provided by a subscription or a provider's API key.

Notification — a message about what's happening in the team (completed tasks, approval requests). Find them under the bell icon in the top bar; important alerts can also go to Telegram.

Onboarding — the setup wizard after your first sign-in: your details, the company profile, and connecting an AI provider. The information you fill in becomes the team's baseline context. See Getting started.

OpenRouter — a service that makes many AI models from different makers available through a single API key. Paid with credit; you see the usage in the **Costs** section. See Costs and limits.

Prompt (assignment) — the text you use to tell an agent what to do. The more specific the assignment (goal, audience, scope, deadline), the better the result.

Review — a task state waiting for assessment. In the detail you either **Approve** it or return it with comments using the **Rework** button. See Tasks.

Scheduled job — a job that runs automatically on a schedule, for example the morning briefing or a weekly report. Managed in the **Task Scheduler**. See Task Scheduler.

Skill — an agent's ability: a playbook and connections that let it do a specific thing (write SEO articles, publish posts, generate images…). See Skills.

Source (wiki) — the original material a wiki page was created from: a meeting transcript, a document, a selected email, or a job's output. Every piece of information in the wiki can be traced back to its source. See Team wiki.

Subagent — a helper an agent spins up for a part of a large assignment. Thanks to subagents, the team handles extensive work in parallel — several parts at once.

Synthesis (wiki) — a summary wiki page that connects information from several sources into one bigger picture (e.g. everything about one project). See Team wiki.

Thread (chat) — a separate conversation with an agent. You can have several threads with one agent (e.g. one for a campaign, another for reporting) — start a new one with **New chat**. See Chatting with agents.

Token (access) — an access key to an external service (also called an API key). It works like a password for an application: you enter it in Integrations and can revoke it at any time.

Transcription (voice) — the automatic conversion of spoken word to text. It works for the microphone in chat and for voice messages in Telegram — the agent then acts on the transcribed text.

Webhook — an address where an external service (e.g. Fireflies or a form) automatically sends a message about an event. Inbound webhooks are managed in **Settings → Automations**.

Wiki — company memory in **Data Center**: it automatically builds tidy pages from selected meetings, documents, and emails, with traceable origins. See Team wiki.