interweb

Generated on 21/10/2026 17:23

KNOWLEDGE BASE · PUBLIC DOCUMENT

What OpenClaw costs: the API keys, not the server

AI & Agents

CODE	VERSION	COUNTRY · LANGUAGE
RLX-MZ-KB-946	v2026.10.10	Mozambique · English

Interweb Mozambique
RLX Holdings, S.A. · Group RLX Holdings
Tax ID (NUIT) 401 191 607 · Legal entity no. (NUEL) 105 045 974 · Economic activity code (CAE) 62022

https://www.meu.interweb.co.mz/knowledgebase/946

What OpenClaw costs: the API keys, not the server

AI & Agents Created on **October 10, 2026** 3 min read



OpenClaw is open source and you do not pay to use it. What you pay for, besides the server, are the **calls to the AI model**: every time the agent thinks, it sends text to a provider, and the provider charges for what it processes. That is the running cost that grows with use, and it is that one, not the VPS, that tends to surprise people.

We do not write model prices here, because they change and differ from provider to provider. The current price is always on the **pricing page of the provider** of the model you choose (OpenAI, Anthropic, Google or another). The server price is on the **VPS** page.

The two separate bills

Bill	Who charges	How it behaves
Server (VPS)	Interweb, on a normal invoice	Predictable: the same whether the agent is busy or idle.
AI model	The model provider, through the API key	Variable: it rises when the agent works more, reads more text or runs scheduled tasks.
Extras (voice, images, web search)	The respective providers	Each connected service has its own bill.

What makes the model bill go up

1. **Long conversations.** As a rule the history is sent to the model again with every request. A long conversation costs more with every reply. See **how tokens are counted**.

2. **Big files and pages.** If the agent reads a whole document to answer a short question, you pay for the whole document.

3 **Scheduled tasks.** A task that runs by itself every hour spends even when you are not looking.

4 **Bigger models.** The more capable models usually cost more per request than the small ones. Use the big one where quality matters and the small one for the rest.

5 **Retry loops.** An agent that fails and retries without stopping can spend a lot before anyone notices.

How to keep the bill under control

1 **Set a limit at the provider.** Most let you define a spending cap or usage alerts in the console. Look for that option before you switch the agent on.

2 **One key per agent.** That way you see what each one spends and you can revoke one without stopping the others.

3 **Check usage after a week.** The provider's console shows use per day. If there is a spike, look for the task that caused it.

4 **Choose the model for the job.** Summarising and sorting rarely need the most expensive model.

! **A leaked key is an open tab.**
Whoever holds it spends on your account. Never put it in code visible on a website or in a public repository. See **how to keep a key safe**.

If you want a more predictable cost, you can run models on your own VPS. In return you need a much bigger machine and more upkeep. See **when it makes sense**.

RECOMMENDED

Need to compare server plans for the agent? See the specifications and prices on the VPS page.

See VPS plans →

SEE ALSO

What OpenClaw is, and what you need to run it →

- **Tokens explained: how AI model APIs are counted** →
- **Getting an API key for an AI model, and keeping it safe** →
- **VPS plans** →

interweb

Interweb Mozambique

RLX Holdings, S.A. · Group RLX Holdings

TAX ID (NUIT)	401 191 607
LEGAL ENTITY NO. (NUEL)	105 045 974
ECONOMIC ACTIVITY CODE (CAE)	62022
REGISTERED OFFICE	Bairro 7 de Setembro Vilankulo 1304 · CEP 0411-02, Inhambane, Mozambique
CONTACT	info@interweb.co.mz · +258 84 635 1111 · +258 86 635 1111 · www.interweb.co.mz
OPENING HOURS	Monday to Friday - 08 AM to 04 PM Saturday & Sunday - Closed

RLX-MZ-KB-946 · v2026.10.10 · Generated on 21/10/2026 17:23

Online version of this document