interweb

Generated on 28/11/2026 22:24

KNOWLEDGE BASE · PUBLIC DOCUMENT

Why Aider shows a cost of 0, and where to see what you really spent

Aider

CODE	VERSION	COUNTRY · LANGUAGE
RLX-MZ-KB-3228	v2026.11.28	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/3228

Why Aider shows a cost of 0, and where to see what you really spent

Aider Created on **November 28, 2026** 2 min read

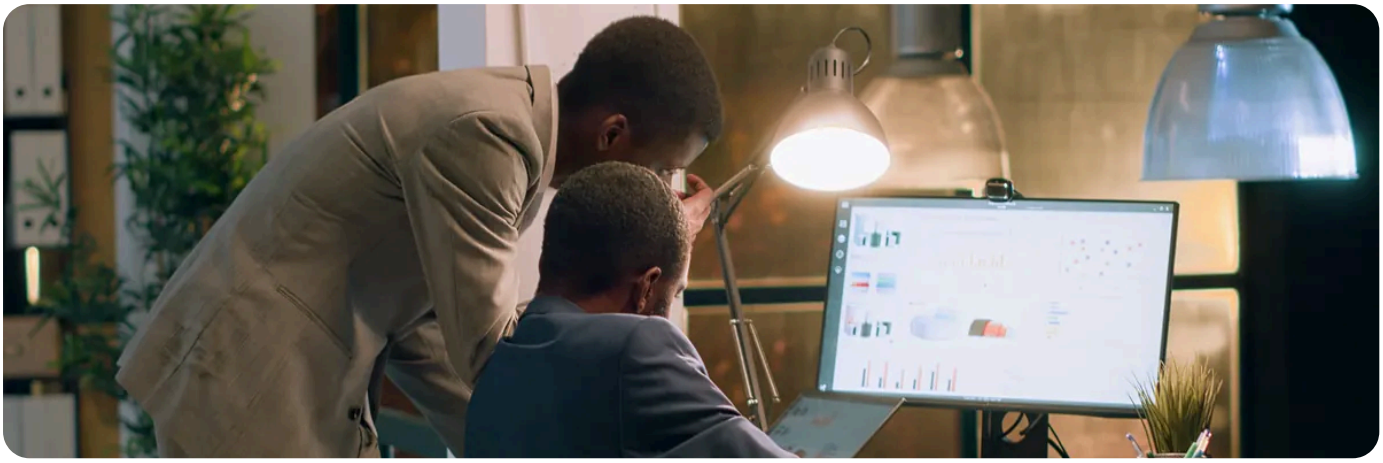

Aider works out the cost of a chat from the data it has on each model. With All in One AI models, that cost stays at 0. It does not mean the requests were free.

Why it stays at 0

Aider does not know our models. Without data, it assumes the model is free, as its **documentation** says. If you created the `.aider.model.metadata.json` from our articles, the `input_cost_per_token` and `output_cost_per_token` fields are 0 on purpose: Aider counts cost per token, and our requests are paid with AI credits or Coding Plan points, which it does not know.

Where to see what you spent

Where	What it shows
History tab	Every AI request made through the API, with the model, the key and the credits spent. You can filter, see the details and export to CSV.
Plan tab	With a Coding Plan, the This week and This month bars, with the points used and the points left.
`aio balance`	The credits and, with a plan, the points left this week and this month.
`GET /v1/ai/balance`	The same, in JSON, for use in a program.

A key just for Aider

In the history, each request shows the key it used. If you create a key just for Aider in the **API Center**, with `ai:read` and `ai:write`, you know at once which requests were its own. If you ever want to turn it off, your other programs keep working.

Knowing the price beforehand

The price of each model is in the **Models tab** and in `GET /v1/models`, in the `pricing` field, in credits per million tokens. Input tokens are paid too, and in Aider `/tokens` shows how many the current chat takes up. Fewer files in the chat means cheaper requests.

How each request is charged, with the reservation and the final adjustment, is in **what a chat request costs through the API**.

RECOMMENDED

Want to know what you spend coding with AI? Visit the Coding with AI page.

See Coding with AI →

SEE ALSO

- **Aider and the Coding Plan: points instead of credits**
- **Checking the points left: the Plan tab, GET /v1/ai/balance and aio balance**
- **What a chat request costs through the API**
- **Aider with All in One AI: the two variables, the openai/ prefix and the context window**
- **Coding with AI in All in One AI**

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

Online version of this document

RLX-MZ-KB-3228 · v2026.11.28 · Generated on 28/11/2026 22:24