interweb

Generated on 28/11/2026 21:58

KNOWLEDGE BASE · PUBLIC DOCUMENT

The minimum of 0.01 credits per request and how costs are rounded up

AI Credits & Costs

CODE	VERSION	COUNTRY · LANGUAGE
RLX-MZ-KB-2974	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/2974

The minimum of 0.01 credits per request and how costs are rounded up

AI Credits & Costs Created on **November 28, 2026** 2 min read



Credits are counted with two decimal places. The cost of each request is rounded up to the nearest hundredth, and no charged request costs less than **0.01 credits**.

The example from the documentation

On a model at 56 credits per million input tokens and 112 per million output tokens, a request with 1,000 input tokens and 500 output tokens costs 0.056 + 0.056 = 0.112 credits. It is rounded up: you pay **0.12**.

Very small requests

On the same model, a question with 20 input tokens and a 10-token answer comes to 0.00112 + 0.00112 = 0.00224 credits. That is below the minimum, so you pay **0.01**.

The minimum matters when you make lots of small requests: 1,000 requests cost at least 10 credits, however short they are.

What costs nothing

Case	Why
Requests with a `sandbox_` key	They do not call the model and charge nothing. `usage` comes back as zero.
A `402` for lack of credits	The request never reached the model. Nothing was charged.
An image, video or voice that fails	The whole reservation goes back to your balance.

Where you see what was charged

In the Credits column of **History**, in the "Real cost" the Playground shows after each answer and in your balance, always with two decimal places.

Lots of tiny tasks, such as sorting 1,000 short sentences? Sending several in the same message avoids paying the minimum 1,000 times. See **how to prepare a large batch**.

RECOMMENDED

Want to see what each request cost? Open History.

See All in One AI →

SEE ALSO

- **What a chat request costs through the API: the reservation, max_tokens and the final adjustment**
- **Reading a model's price: credits per million tokens, per image, per second or per generation**
- **Before a large batch of AI requests: checking the balance and the price first**
- **The History tab: what each AI request spent, with filters and CSV export**
- **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-2974 · v2026.11.28 · Generated on 28/11/2026 21:58