




RESPONSES PANEL

Requires an AI subscription

 **What you can do here:** Review all your AI interactions and usage. See a complete log of AI requests, responses, costs, and performance metrics. Useful for tracking your AI usage, reviewing past queries, and understanding costs. Great for administrators and power users. This panel shows a searchable table of AI calls (rows from the `ai_logs` table). You must be logged in to see your logs.

What kinds of AI work get logged (Type column):

- **ai-coding:** “AI coding” runs that read source text (often in chunks) and generate things like coded links / structured outputs.
- **vignette:whole / vignette:typical:** Vignette generation (AI-written narrative summaries).
- **answers:question_expansion:** AI Answers expanding your question into multiple search phrases.
- **answers:answer:** AI Answers generating an answer from retrieved evidence.
- **filter:soft-recode:generate_magnets:** Filter Pipeline (Soft Recode / SRP) using AI to propose magnet labels.
- **filter:soft-recode:label_cluster_node:** Filter Pipeline (Soft Recode / SRP) using AI to label individual cluster nodes.
- **filter:cluster:label_clusters:** Filter Pipeline (cluster filter) using AI to label clusters.
- **docs:chat:** MapCat documentation chat. Logged but **not counted** toward your monthly credits.
- **– (blank):** Older log rows (before we added `request_type`) or any legacy code path that inserts into `ai_logs` without setting a type. You can still click the eye to see the full prompt/response.

How to use the table:

- **Scope toggle (top-right):**
- **This project:** shows only AI calls from the currently selected project (exact match).
- **All projects:** shows AI calls across all projects you have access to.
- **Sort:** click a column header (e.g. **Timestamp**, **Cost**) to sort by that column; click again to reverse the sort.
- **Filter row (the grey row under the headers):** type/select a value and the table refreshes automatically.
- **Date:** pick a date. If you pick a past date, it filters **from that date up to today**. If you pick today, it filters **today only**.
- **Project:** free-text filter (substring match). Most useful when Scope = **All projects**.
- **Model:** dropdown.

- **Status: Success or Error.**
- **Source:** filter by `source_id`.
- **Prompt:** filter by text inside the prompt (substring match).
- **Clear filters (X button):** appears only when any filters are active; click it to reset the filters.
- **Pagination:** use the pager at the bottom to move through pages; use the **page size** dropdown to change rows per page.

What the columns mean (quick guide):

- **Timestamp:** when the AI call was made.
- **Project:** which project it was for.
- **Model / Region:** which model ran and which region it used (where applicable).
- **Status:** success / error (some rows may show “pending”).
- **Time (s):** response time.
- **Chunk / Iter / Batch / Source:** metadata used mainly for chunked/batched runs.
- **Prompt:** a short preview (hover to see more).
- **P Tok / C Tok / Cost:** prompt tokens, completion tokens. **Admins** see cost in USD (\$). **Non-admins** see **Credits** (1 credit \approx \$0.01; users without paid AI plans get 10 free credits per month; credits renew monthly and do not roll over).
- **User:** the user email associated with the call (when available).
- **Type:** what feature created the call (see “What kinds of AI work...” above).

Credits and limits: Non-admins see **Credits this month: X / Y** in the summary (usage / limit).

Credits renew at the start of each month and do not roll over. MapCat ([docs:chat](#)) is logged but excluded from the usage count. If you exceed your monthly limit, a warning appears and AI features are disabled until next month.

Click the eye (Details column) to see the full record:

- Shows **full prompt**, the **full response** (and sometimes renders it as a table if it looks like JSON/tabular text), plus **error messages** and the **raw JSON response** for debugging.