

Ask (global)

ask-global.html · 1360×820

ChatMaestro 1/2026.08 2 PN priya.n

Ask 3 Global · all contexts Read-only

4 Search saved questions...

5 Conversation No scope

PN Which experiments' runs cost the most per composite point, and how many of each used a local judge?

Three experiments stand out. **Politeness · SLM** is the most expensive at **\$0.041** per composite point — its runs lean on a hosted judge. **Terse-prompt** is the most efficient (**\$0.028**), and **2 of its 4 runs** used a local judge. Local-judge runs averaged **31% cheaper** at similar scores.

Experiment	Runs	\$ / composite	Local-judge runs
Politeness · SLM	5	\$0.041	1 of 5
Nudge-B sweep	6	\$0.034	3 of 6
Terse-prompt	4	\$0.028	2 of 4

Local vs hosted judge — avg run cost

read-only · validated Show query · joins 4 tables 6 Save prompt

Ask anything across the workspace... Ask

Annotations

- Version pill (right-aligned)** Dimmed build pill sits on the RIGHT, beside the account menu. Hover reveals the build sha; on phones it folds into the account menu.
- Account menu** Avatar + name + chevron open the account menu (profile, settings, sign out). On narrow screens the feedback / bug / dark-mode icons and version pill collapse into here.
- No scope** The global Ask is the same component with no `scope_key`. Because there's no scope, a question can **cross entities** — this one joins experiments, runs, results and model configs in one read-only query. Ask across everything at once — experiments, runs, results and models.
- Saved-question catalog** Browses **every** saved question you may see, grouped by context and filterable by visibility (mine · approved · built-in). A retired question is not here: it stays on the Saved-questions screen for admins but leaves every Ask surface until restored. Ready-made questions you can run, grouped by area.
- Read-only** — questions run with your own access, and nothing gets changed.
- Save a question** Save prompt files a new saved question here, just as the contextual Ask does.