

Dashboard

dashboard.html · 1360×820

WORK

Dashboard

Experiments

Nudges

Model Configs

Cohorts

MONITOR

Live Monitor

Results

Ask

Saved Questions

COMMS

Messages

Documents

OVERSIGHT

Audit Trail

Issue Tracker

ADMIN

Experimenters

Enrollees

Model Catalog

ChatMaestro

2026.08

PN priya.n

3

View: ExperimenterAdmin

Dashboard · Your experiments at a glance

3Running now

9Experiments

12Active enrollees

5Nudges

\$248Spend MTD

Active Runs3 live

Terse-prompt studyPilot-20 · nudge A

● RUNNING7watch

Politeness · SLM ensemblePilot-20 · polite A

● RUNNING20watch

Redundant-prompt · Opus · hotGroup-C · hot A

● PAUSED18watch

Self-consistency probeBatch-7 · run 3

● RUNNING22watch

Quick Actions

+ New experiment

New model config

New nudgeNew cohort

Recent Resultslast 4 finished

Terse-prompt studyPilot-20 · nudge B · Jul 16

● 71.2+2.7

Self-consistency probeBatch-7 · run 2 · Jul 15

72.8+1.4

Judge-calibration A/BPanel-A · Jul 13

70.1— 0.0

Verbosity-bias checkPilot-20 · Jul 11

64.7-3.1

Spend to Date

\$248.60this month

Opus 4.5\$142.08

Sonnet 4.5\$74.30

Haiku 4.5\$32.22

Annotations

1

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.

2

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.

3

Role-aware view switch The Dashboard is one home screen with two lenses. ‘Experimenter’ (default) shows priya.n’s OWN active runs, results and spend. ‘Admin’ swaps in an org-wide experimenter-activity overview — visible because priya.n also holds the admin role.

4

At-a-glance KPI tiles A compact KPI strip summarizing priya.n’s workspace: running now, total experiments authored, enrollees across live runs, and month-to-date spend. Each tile is a quiet card — number + label — the numbers reconcile with the cards below.

5

Active runs (RUNNING / PAUSED) Only runs with live state land here — RUNNING (teal) or PAUSED (orange). ‘Active’ is a derived condition of the run, never a property of the reusable experiment package. Each row jumps straight to the Live monitor via ‘watch’.

6

RUNNING — the run is actively streaming chat rounds; enrollees are answering live.

7

Jumps to the Live monitor filtered to this run.

Copyright © 2026 Agilis Software LLC

Page 1 of 2

-
- 8 Composite = weighted rubric blend** Composite (0–100) is a weighted blend of the run's rubric factors, reported per finished run. The trend hint compares against the run launched most recently before this one **in the same experiment**, and shows nothing when there is no predecessor. It is scoped that way because a composite is only comparable inside one experiment, which fixes the rubric, the judge and the weights — this list spans experiments, so an arrow drawn across two of them would compare different measuring instruments. It's a directional cue, not a significance test — the Results screen carries the full breakdown.
-
- 9** Composite 71.2 — weighted blend of the rubric's quality factors for this run.
-
- 10 Spend accrues per model** Running total of LLM inference cost charged to priya.n's runs, broken down by the model that produced the tokens. Owner holds the only provider keys, so this is the experimenter's read-only view of what her runs have consumed.