

Mercury 2.5

Inception · detected 2026-09-09 06:41 UTC · alert out 5772s later · 14 alerts delivered

AT A GLANCE

PRICE PER 1M

\$0.04 / \$0.15

input / output · cached
\$0.004

CONTEXT

260K

tokens in

MAX OUTPUT

66K

tokens per reply

TAKES


text

returns text


HOW IT COMPARES

MODEL	ROLE	IN / 1M	OUT / 1M	CONTEXT	BENCH INDEX
Mercury 2.5	new	\$0.04	\$0.15	260K	not rated yet
Mercury 2	it follows	—	—	—	—
Claude Opus 5.5	today's #1	\$4	\$20	1M	100
Claude Sonnet 5.5	best value	\$2	\$10	1M	93.3
vs Claude Opus 5.5	99% cheaper today's #1: \$4 / \$20, bench index 100; context 1M				
vs Claude Sonnet 5.5	input 98% cheaper, output 99% cheaper best value: \$2 / \$10, bench index 93.3; context 1M				

PRICE PER 1M TOKENS, BLENDED 3 IN : 1 OUT



WHERE THE PUBLIC LEADERBOARDS STAND TODAY



Mercury 2.5 is not rated yet: the public boards usually score a new model within days, and its page on artificialwatch.com gains the scores then. The bench index is the median of each model's percentile across the public leaderboards that rate it (LMarena, LiveBench, SimpleBench, LLM Stats, Vals AI, Artificial Analysis).

Sources: price, limits and modalities for Mercury 2.5 as the OpenRouter catalog lists inception/mercury-2.5; the other models as their makers list them (models.dev); rankings from the ArtificialWatch Bench Index. Facts only: no forecast of ours, no claims from the lab's announcement.