LongCat 2.5

Meituan

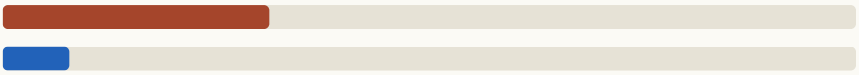
AT A GLANCE

PRICE PER 1M	CONTEXT	MAX OUTPUT	TAKES
not listed	– tokens in	– tokens per reply	**text** returns text

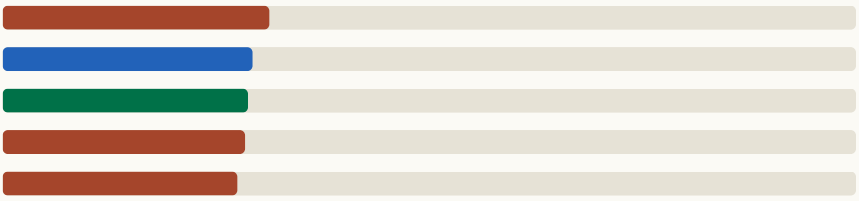
HOW IT COMPARES

MODEL	ROLE	IN / 1M	OUT / 1M	CONTEXT	BENCH INDEX
LongCat 2.5	new	–	–	–	not rated yet
Claude Fable 5.2	today's #1	$8	$40	1M	100
Gemini 4 Carbon	best value	$2	$10	1.05M	93.7

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



WHERE THE PUBLIC LEADERBOARDS STAND TODAY



LongCat 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, Epoch AI).

Sources: price, limits for LongCat 2.5 as the OpenRouter catalog lists it; 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.